

Regulated Revenues and Hospital Behavior: Evidence from a Medicare Overhaul*

Tal Gross[†] Adam Sacarny[‡] Maggie Shi[§] David Silver[¶]

August 2022

Abstract

We study a 2008 policy reform in which Medicare revised its hospital payment system to better reflect patients' severity of illness. We construct a simulated instrument that predicts a hospital's policy-induced change in reimbursement using pre-reform patients and post-reform rules. The reform led to large persistent changes in Medicare payment rates across hospitals. Hospitals that faced larger gains in Medicare reimbursement increased the volume of Medicare patients they treated. The estimates imply a volume elasticity of 1.2. To accommodate greater volume, hospitals increased nurse employment, but also lowered length of stay.

*We thank Caitlin Carroll, Josh Gottlieb, Peter Hull, workshop participants at the AEA/ASSA ASHEcon session, the ASHEcon virtual conference, Whistler Health Economics Summit, Columbia University, Princeton University, UC Berkeley, UCSB, and the Congressional Budget Office for their helpful feedback.

[†]Questrom School of Business, Boston University and NBER.

[‡]Mailman School of Public Health, Columbia University and NBER.

[§]Department of Economics, Columbia University.

[¶]Department of Economics, Princeton University and NBER.

In the United States, the federal government pays private hospitals through the Medicare program, the largest insurer in the country. For decades, Medicare has paid hospitals a lump-sum fee for each visit. In 2019, those payments alone represented 0.9 percent of Gross Domestic Product (MEDPAC, 2021). In turn, the level of per-visit payments is a key tool of healthcare policymakers seeking to cut costs. It is also a constant subject of public debate.

Changes to that per-visit reimbursement rate are not simple exercises in arithmetic. Providers can react to reimbursement cuts by adjusting the number of patients they treat. On the one hand, lower payment rates would seem to give providers an incentive to see fewer patients, compounding the savings from rate cuts. On the other hand, if healthcare providers exhibit “wealth effects,” they could respond to lower reimbursement rates by increasing their patient volume. Such a response would limit the reimbursement cuts’ effectiveness.

The magnitude of this volume response on the part of providers is thus a key input into policy debates over the future of fee-for-service reimbursement. A full understanding of hospitals’ volume responses is especially important given that a new generation of hospital payment reforms penalize or reward hospitals by modifying per-visit Medicare reimbursements. Through a number of programs, Medicare adjusts its reimbursements to incentivize lower readmission rates, improved survival, more evidence-based care, and fewer hospital-acquired infections.¹ All of these programs change per-visit payment rates based on indicators of performance.

This paper describes how hospitals respond to such changes in payment rates. In 2008, Medicare permanently changed prospective reimbursement rates for hospitals, redistributing roughly \$2 billion across hospitals annually in the process. That amounts to roughly four times the annual penalties imposed by Medicare’s value-based program to reduce readmissions and seven times the annual penalties imposed by its program to reduce hospital-acquired conditions. Similar to the payment-rate effects of value-based programs, some hospitals, by virtue of their patient mix, enjoyed a highly persistent relative increase in Medicare reimbursements and other hospitals saw a relative decline.

We study how these changes in reimbursement rates affected hospitals’ behavior using a “simulated-instruments” approach (Currie and Gruber, 1996a,b). The results demonstrate—first and foremost—that hospital care follows a traditional, upward-sloping supply curve. Increases in Medicare reimbursement rates led hospitals to expand their volume of Medicare admissions. Our point estimates imply a hospital-level volume elasticity of 1.2. To operationalize this result, we note that CMS’s value-based program to reduce hospital readmissions imposes a 3-percent penalty on the lowest-performing hospitals. Our findings imply the program would reduce hospitalizations by approximately 4 percent through supply responses

¹Appendix A summarizes those programs.

alone.

We document negative effects of payment rates on length of stay, and positive effects on other inputs, most notably nurse employment. Such a pattern is consistent with hospitals responding to higher rates by increasing their overall throughput. It is more difficult for us to estimate an effect on quality of care, because the composition of patients in the hospital changed. That said, the results suggest that, if anything, higher payment rates raised quality of care. Viewed in light of these findings, value-based programs' rate penalties likely save money but, in doing so, may impede access.

These effects involve only Medicare patients, but we also find that volume effects “spill over” to other patients covered by Medicaid. And yet, Medicaid reimbursements were not intentionally targeted by the reform. There are several channels through which this spillover might operate, including organizational frictions that prevent hospitals from tailoring their responses to insurers (Frandsen et al., 2019), mimicking of Medicare's reforms by other insurers (Clemens and Gottlieb, 2017), and hospitals using Medicare revenue to finance care for other patients (Frakt, 2011). Regardless of which mechanism is most relevant, these results highlight that one payer's payment is closely linked to the care that other payers' patients receive.

This paper contributes to a broad literature examining how healthcare providers respond to reimbursement. One clear finding is that hospitals strategically change how they bill Medicare for patients when certain diagnoses become more lucrative, a phenomenon often called “upcoding” (Silverman and Skinner, 2004; Dafny, 2005; Sacarny, 2018). The remainder of the literature has largely focused on “real” (as opposed to coding) responses. Cutler (1995) examines the initial shift to the PPS from retrospective hospital payment. Studies of physician behavior have generally found positive volume responses to payment rates (Clemens and Gottlieb, 2014; Alexander and Schnell, 2019; Baker and Royalty, 2000; Currie et al., 1994). Clemens and Gottlieb (2014) estimate a price elasticity of supply of outpatient services of 1.4, driven by increased intensity of care per visit rather than increased volume. This mirrors our findings and is consistent with the different incentive systems underlying Medicare's payments to physicians (fee for service) versus hospitals (prospective payment): higher fee-for-service prices encourage more services per visit, while higher prospective payments encourage hospitals to treat more patients, perhaps by reducing intensity or time per patient.

In the hospital context, there is a growing body of work investigating supply-side responses to financing changes (e.g. Duggan, 2000; Baicker and Staiger, 2005; Acemoglu and Finkelstein, 2008; Kaestner and Guardado, 2008; Wu and Shen, 2014; Chang and Jacobson, 2011; Cooper et al., 2017; Azoulay et al., 2020). Acemoglu and Finkelstein (2008) study

the introduction of the hospital prospective payment system in 1983, finding that hospitals reacted by cutting labor inputs and shifting towards capital. More recently, Cooper et al. (2017) consider the political determinants of hospital payment and find that politically induced rate increases lead to increased employment and investment, while Azoulay et al. (2020) find that rate *cuts* stemming from the 1997 Balanced Budget Act led academic medical centers to substitute towards research.

This paper contributes to the literature by providing novel estimates of hospital supply responses to one of Medicare’s primary policy levers: reimbursement rates. Medicare reimbursements amount to roughly half of hospital revenue, and there exists an ongoing public debate regarding how Medicare reimbursement to hospitals ought to change. Much of this debate focuses on scope and intensity of value-based payment reforms, which tend to manipulate hospital prices to incentivize better patient outcomes. The extent to which hospitals react to price shocks by admitting more or fewer patients thus plays a crucial role in understanding the net effects of these pervasive policies.

1 Data, Institutional Background and Empirical Strategy

1.1 Data

This study’s main data source is the MEDPAR file, which contains records of all Medicare inpatients from 2002 through 2013. We rely on that file to calculate the reimbursement shock for each hospital as well as for the key outcomes of hospital volume and average length of stay. We also use that file to observe realized and price-adjusted hospital payments. The data contain all of the attributes of each hospital stay that determine the Diagnosis-Related Group (DRG), such as diagnoses and procedures, which we feed into DRG software to construct which DRG each patient would have counterfactually fallen into in other years. We also use Medicare data on enrollment, disease histories, and outpatient hospital visits.

To observe hospital scale and labor inputs, we draw on the American Hospital Association (AHA) annual surveys and Centers for Medicare and Medicaid Services (CMS) Provider of Services files from 2003 through 2013. Additional information on payment comes from public CMS rulemaking data published annually in the Federal Register.

1.2 Institutional Background

CMS has prospectively reimbursed hospitals for Medicare-insured admissions since the inception of Inpatient Prospective Payment System (IPPS) in 1983 (Cutler and Zeckhauser, 2000). Under this system, each case is categorized into a base DRG, such as heart failure,

based on the diagnosis or procedure that constituted the focus of treatment. The base DRG is further categorized into DRGs, the unit of payment, by adjusting for patient severity differences. Hospitals are typically reimbursed a predefined amount for the entirety of that case's treatment, with differences in payment stemming in large part from what is known as the DRG weight, a unitless number that indicates the relative price CMS attaches to that DRG.

The DRG system that existed prior to the MS-DRG reform recognized severity differences within base DRGs via a list of complications and comorbidities (CCs) that triggered higher reimbursement. The pre-reform list of CCs relied heavily on patients' histories of relevant chronic diseases. This list had lost its power to discriminate hospital resource-use over the first 22 years of IPPS. Strategic coding of these CCs, alongside compositional changes in hospitalizations, led to the overwhelming majority (78 percent) of admissions having a CC on the eve of the reform (CMS, 2006 pp. 47154).

Recognizing this deficiency, policymakers set out to overhaul the DRG system to "better recognize severity of illness" (Centers for Medicare and Medicaid Services, 2007). The resulting switch to Medicare Severity DRGs (MS-DRGs), which was phased in from 2007 to 2009, marked the most significant overhaul of CMS's reimbursement system since the introduction of prospective payment. The new MS-DRG system introduced completely new lists of qualifying CCs, alongside a third, highest tier of reimbursement within many base DRGs for patients with "major" CCs (MCCs). In the new system, CMS also re-focused CCs and MCCs focused on acute exacerbations of illness rather than histories of chronic disease, since acute illnesses were more predictive of costs of care (CMS, 2019 provides a detailed description of the development of MS-DRGs).

Hospitalizations for stroke provide a useful illustration of how the MS-DRG system altered reimbursement rates depending on case severity. Appendix Figure A3 plots the weights for stroke hospitalizations between 2000 and 2013. Prior to the reform, all stroke cases were reimbursed equally with a DRG weight of 1.2, about \$11,000 per case in 2006. The reform split this base DRG into three tiers: without a CC, with a CC, or with an MCC. By the time the MS-DRG system was fully phased in, in 2009, stroke hospitalizations with an MCC were reimbursed at 219 percent of the rate for those without a CC, and at 157 percent of the rate for those with a CC. The new CC and MCC subcategorizations of base DRGs – stroke as well as many others – substantially expanded the set of total DRGs from 538 to 745.

This led to large shifts in effective reimbursement rates across hospitals because hospitals differed in the shares of their patients with newly defined CCs and MCCs. Patients with MCCs comprised just 14 percent of all admissions at the 10th percentile hospital, but 29 percent of all admissions at the 90th percentile hospital (Figure 1a). Those differences form

the basis for our empirical strategy.

1.3 Empirical Strategy

Our study seeks to isolate the predictable mechanical effects of the switch to MS-DRGs on hospital reimbursements. We then relate these effects to changes in hospital behavior to understand how hospitals respond to reimbursement rates.

The pre-reform differences in shares of admissions that would have been categorized with an MCC had they occurred after the reform’s implementation (Figure 1a) provide a natural starting point for quantifying the differential impacts of the reform across hospitals. As these differences are calculated using admissions that occurred in 2006, before the phase-in of the reform, they reflect pre-existing patient-mix differences that were not contemporaneously factored into reimbursements. As hospitals’ patient mixes tend to be relatively stable over time, hospitals with higher shares of admissions with MCCs in 2006 stood to see larger increases in their reimbursement rates once the reform came into effect.

To more formally characterize the predictable mechanical effect of the MS-DRG reform on hospitals’ reimbursement rates, we generalize the insight above. Our approach is borrowed from the “simulated instruments” literature (e.g. Currie and Gruber, 1996a,b), adapted to our setting. We continue to focus exclusively on admissions from 2006.

Our approach requires two key ingredients for each admission (indexed by i) observed in 2006: (a) its actual DRG weight that was used for payment in 2006, w_i , and (b) its counterfactual MS-DRG weight that would have applied had the admission occurred in 2009, after the reform had fully phased in, w_i^* . The first is observed in our data. The second we simulate by reprocessing admissions under the 2009 MS-DRG rules. This reprocessing step uses each admission’s listed diagnosis and procedure codes to assign an MS-DRG to each admission (including CCs and MCCs) and applies the relevant weight from the 2009 payment system. The difference between the counterfactual and actual weights ($w_i^* - w_i$) provides the simulated admission-specific difference in payment rates induced by the reform.

To capture the hospital-level impact of the reform, we average these differences at the level of the hospital at which the admission occurred. We call this difference hospital h ’s “reimbursement shock”:

$$\text{SHOCK}_h = \frac{1}{N_h} \sum_{\{i:h(i)=h\}} (w_i^* - w_i), \quad (1)$$

where $h(i)$ indicates the hospital at which admission i occurred and N_h is the count of the admissions at hospital h in 2006. Since both terms in this difference are computed using

the same admissions cohort receiving identical diagnoses and care, this measure captures the change in reimbursements hospital h would have experienced between 2006 and 2009 had the hospital changed nothing about its operations, admissions, or coding.

These measured revenue shocks are substantial for many hospitals. Figure 1b presents a histogram of this reimbursement shock across hospitals. The distribution is centered on zero: on average, hospitals did not face a fundamental change in their reimbursement rates, based on pre-reform patient mix. Indeed, the reform was designed to hold this shock to zero at the national level.

But at the hospital level, there exists a great deal of mass beyond zero. The gap between the 90th and the 10th percentile hospital represents a 7.1-percent difference in average Medicare reimbursement rates resulting from the policy change. Given an average Medicare reimbursement in 2006 of \$8,618, this translates to an increase in per-admission reimbursements of \$612. Further, given an average caseload of 4,123 discharges, this translates to a \$2.5 million change in hospital-level Medicare revenue per year. Overall, the reform redistributed almost 2 percent of inpatient revenue, or \$1.95 billion, across hospitals. Since 49 percent of the revenue went to hospital gains and 51 percent went to hospital losses, the predicted net effect of the reform on total reimbursements was -0.05 percent, strikingly close to zero.²

Unsurprisingly, the reimbursement shocks are closely related to hospitals' pre-reform MCC shares. The third panel of Figure 1 displays the empirical relationship between $SHOCK_h$ on the vertical axis and a hospital's share of cases with MCCs on the horizontal axis. The two are indeed highly correlated ($\rho = 0.69$), with residual variation in $SHOCK_h$ deriving from differences in shares of patients without any CCs, as well as differences in the types of diseases that hospitals treat, some of which were more impacted by the reform than others.

Variation in the revenue shock across hospitals could be driven by pre-existing differences in patient health status, differences in how intensively hospitals code the diagnoses of patients, or both. As in most settings that involve claims data, it is difficult to distinguish these sources of variation. The 2008 reform caused certain patient populations to become more lucrative than they previously had been, since those patients became categorized as having MCCs. To the extent that there are variations across hospitals in the true shares of their patient populations with MCCs, our approach will predict differences in their revenue shocks. It is also possible that some variation in the revenue shock was driven by hospitals' coding practices. That said, the relationship between hospitals' revenue shocks (i.e. our in-

²We calculate the revenue redistributed by taking the absolute value of the predicted shock based on a hospital's 2006 patients and multiplying it by the hospital's 2006 patients and its 2006 conversion factor from DRG weight to dollars.

strument) and actual Medicare reimbursement rates was close to unity and persistent. That suggests that variations in the shock were not the result of some hospitals focusing their coding intensity on the payment system in place at the time – such variations would not persist post-reform as these hospitals would come to code more intensively according to the new rules.

The reimbursement shock captures several ways in which the policy change affected the incentives faced by hospitals. First, it captures a kind of wealth effect. If hospitals were to change nothing about the patients they admit and the way they code those patients' conditions, hospitals with high values of the reimbursement shock would enjoy increased payments per patient.

Second, the reimbursement shock describes the increased profitability of the patient pool *ex ante*. Those patients, in turn, are plausibly representative of the “potential patients” that the hospital could attract to increase volume because they would draw them from the same catchment areas. In that sense, the reimbursement shock captures the degree to which the policy change created an incentive for some hospitals to increase Medicare volume.

Crucially, the reimbursement shock does *not* measure the actual response of each hospital to the policy change. Some hospitals may have responded by more aggressively upcoding their patients, that is, submitting CCs and MCCs to CMS. That impact is not captured by the reimbursement shock, making the reimbursement shock a plausibly exogenous instrument by which one can study how the reform affected hospitals.

Table 1 explores how this change in average reimbursement rates was correlated with hospitals' characteristics.³ To do so, the table focuses on hospitals in two groups: those with a shock above zero that were predicted to experience an increase in reimbursement, and those with a shock below zero that were predicted to experience a reduction. Hospitals that saw increases were larger and more likely to be non-profit. Those differences in hospital characteristics, however, do not invalidate the research design. As we describe below, this paper's identification strategy amounts to a difference-in-difference comparison around the policy change.

To further explore variation in the revenue shock, Appendix Figure A1 plots average DRG weights over time for the hospitals most affected by the policy change. In particular, the figure compares hospitals with the top decile of the revenue shock—those set to experience the greatest gains—versus hospitals with the bottom decile of the revenue shock—those set to experience the biggest losses. Before the reform, the average DRG weights across those two groups of hospitals were similar. Importantly, the *trend* in average DRG weights was similar before the policy change. But following the implementation of the reform, which

³Appendix Figure A2 describes geographic variation in the revenue shock.

began in 2007, the two time series diverge. Hospitals facing the largest positive revenue shocks saw an average DRG weight of 1.6 by 2012, versus only roughly 1.5 for hospitals facing the most-negative revenue shocks.

In order to study the effects of this variation on hospitals' behavior, we take a continuous difference-in-difference approach. In particular, we explore regressions of the form:

$$y_{ht} = \alpha_h + \alpha_t + \sum_{s=2004}^{2013} \beta_s \times I\{s = t\} \times \text{SHOCK}_h + X'_{ht}\Gamma + \varepsilon_{ht}. \quad (2)$$

This regression studies outcomes, y_{ht} , for hospital h in year t . The variable SHOCK_h is the hospital-specific revenue shock, defined by equation (1). We include hospital-specific fixed effects, α_h , and year-specific fixed effects, α_t . The key coefficients of interest are the β_s 's, which measure the effect of the hospital-specific change in DRG weights both before and after the policy was implemented.

The vector X'_{ht} represents additional time-varying controls. In our main specification it consists of interactions between hospital size in ventiles (defined based on the number of beds in 2005) and the year t . We use these controls to account for a Medicare policy change that allowed low-volume hospitals to request payment rate increases starting in late 2010 (the policy was enacted in the Affordable Care Act; CMS, 2010 pp. 50238 describes its implementation). The bed-size ventiles are highly predictive of whether hospitals took up the low-volume adjustment.

A disadvantage of equation (2) is that it provides a point estimate for each year, rather than a single summary of the effects of the revenue shock. For that reason, we also run regressions of the following form:

$$y_{ht} = \alpha_h + \alpha_t + \beta_{\text{interim}} \times \text{SHOCK}_h \times I\{2007 \leq t \leq 2008\}_t + \beta_{\text{post}} \times \text{SHOCK}_h \times I\{t \geq 2009\}_t + X'_{ht}\Gamma + \varepsilon_{ht}. \quad (3)$$

Equation (3) has the advantage of summarizing the overall impact of the revenue shock with one point estimate, β_{post} . To avoid including short-run implementation effects, we separate out 2007 through 2008, the years of partial phase-in of the reform.

Figure 2 describes how hospitals' revenue shocks affected the Medicare payments they received. The first panel plots estimates of the β_s 's from regression (2) with the outcome being hospitals' actual average DRG weights each year. The figure suggests no relationship between the revenue shock and hospitals' average DRG weights before the reform and then an increase in the predictive power of the shock instrument in 2007 and 2008. That increase in 2007 and 2008 is consistent with the phased implementation of the MS-DRG system during

those years. The figure then suggests a relatively stable effect from 2009 through 2013.

The second panel of Figure 2 presents a non-parametric approach to estimate the same relationship: a binned scatterplot. That panel compares actual changes in hospitals' reimbursement to their revenue shock. The vertical axis plots changes in average DRG weights before and after the policy; the horizontal axis plots the instrument. The graph suggests a roughly linear relationship, suggesting that the linear functional form in (2) is a reasonable representation of the underlying data.

The third and fourth panels of Figure 2 perform the same analysis when the outcome of interest is hospitals' average Medicare payments, in thousands of dollars. Those panels present what could be called this empirical strategy's first stage – below we treat this outcome as the endogenous variable in two-stage least squares regressions. The figures suggest that hospitals' revenue shocks increased their average Medicare reimbursement, and that the underlying functional form of that relationship is approximately linear.

Finally, in order to examine the magnitude of the estimates presented graphically in Figure 2, the first two panels of Table 2 presents the analogous difference-in-difference estimates based on equation (3). The first panel shows that hospital revenue shocks passed nearly fully through to actual DRG weights. Hospitals with a one-unit higher shock see DRG weights rise by 0.97 units. The second panel shows the translation of this shock to dollars of hospital reimbursement per patient. A one-unit higher shock leads to a \$7,134 increase in average total payments, which also include reinsurance payments for extraordinarily high-cost patients and certain payments for medical education and capital costs.

2 Results

We begin by exploring how hospitals responded to their revenue shock by manipulating their patient throughput. We first estimate equation (2) to test how the reform affected the volume of Medicare admissions. The first panel of Figure 3 presents estimates of that event-study regression when the outcome of interest is the logarithm of the number of each hospital's Medicare patients each year. The figure suggests no differential trend before the reform went into effect, consistent with the parallel-trends assumption, and then a clear, steadily rising effect after its introduction.

In order to probe the validity of that finding, the second panel of Figure 3 presents a binned scatterplot that studies the same relationship. For all hospitals, we calculate the change in total Medicare patient volume between 2006 and 2009, and examine that change in volume across quantiles of the revenue shock. The figure suggests a positive, linear relationship: hospitals with larger revenue shocks saw larger increases in Medicare volume.

The linear pattern supports our main specification in which outcomes depend linearly on the revenue shock.

To explore how hospitals achieved this increase in patient volume, we study the average length of stay of Medicare patients. The third panel of Figure 3 presents event-study estimates where the outcome of interest is the logarithm of the average length of stay in days of the hospital's Medicare patients. The figure suggests a clear reversal of trend in 2007, immediately as the policy change went into effect. The final panel of Figure 3 plots the corresponding binned scatterplot. That panel supports the assumption of linearity for this outcome and validates the overall effect estimated via the event-study specification.⁴

The second panel of Table 2 presents reduced-form and 2SLS estimates of the effect of the revenue shock on outcomes related to volume and throughput. As is consistent with Figure 3, the first two rows of the panel suggest that the revenue shock raises patient volume and shortens length of stay by economically meaningful and statistically significant magnitudes. In other words, hospitals increased their throughput, treating more patients but keeping each for less time in the hospital.

To interpret the point estimates of that reduced-form effect, consider a hospital with a projected shock of 0.05, putting the hospital at the 90th percentile of the distribution of revenue shocks. Such a shock represents an expected 4-percent rise in Medicare payments for the average hospital. The point estimates suggest that the hospital's behavior would change: its Medicare admissions would rise by 4.7 percent and its average patient would stay at the hospital for 1.6 percent fewer days. These results demonstrate a large, statistically significant effect of the revenue shocks induced by the transition to MS-DRGs on hospitals' Medicare-covered patient throughput.

To further characterize these magnitudes, the table also presents instrumental-variables estimates. These effects can be interpreted as the effect of raising a hospital's Medicare payment by \$1,000 per patient; according to the first stage, such a payment change would result from a predicted revenue shock of 0.14 units. The estimates show that each thousand-dollar increase in average Medicare payments leads to a 13-percent increase in Medicare admissions and a 4.4-percent drop in average length of stay for Medicare patients.⁵ With an

⁴The drop in length of stay could be driven by both a change in hospital behavior and a change in the composition of patients admitted to the hospital. We explore that dynamic in Appendix Figure A4 by predicting length of stay using patients' race, age, gender, and medical diagnosis. Revenue shocks decreased predicted length of stay, albeit by much less than actual length of stay (Panels a and b). The effect on actual length of stay is also qualitatively similar and, as expected, slightly smaller when we control for predicted length of stay (Panel c). The results suggest that there was at least some change in composition, but most of the actual change in length of stay cannot be explained by patient observables.

⁵Appendix Table A1 stratifies this analysis by hospital ownership. The table suggests no statistically significant differences in volume responses by hospital ownership. In contrast, for-profit hospitals seem to exhibit the largest changes in length of stay.

average pre-reform payment rate of \$8,618, the implied payment-rate elasticities of volume and length of stay are 1.20 and -0.38 .

Appendix Tables A2 and A3 explore the robustness of the regressions for volume and length of stay, respectively. The tables address several sources of confounding variation. For instance, enrollment in Medicare Advantage was expanding during this time period, and expanding differently across counties. That expansion of Medicare Advantage surely affected hospital revenues. And so it is reassuring that the main coefficients of interest are not materially changed when we control for county-specific Medicare Advantage penetration. Similarly, it is unclear which year should be the baseline for the revenue shock instrument—above, we use 2006 as the pre-period, but these tables present results with 2004 or 2005 as alternative baseline years. The tables suggest relatively similar estimates across these alternative specifications.

The results above demonstrate that increases in Medicare reimbursement lead to increases in hospitals' Medicare volume. It is natural to inquire as to the source of that increased volume. The marginal Medicare patients may be drawn from a variety of sources. Some patients may have otherwise been treated in outpatient settings, and still others may have not been treated at all. Still other patients may be a matter of “business stealing:” hospitals with positive revenue shocks attract patients from other hospitals.

Unfortunately, the nature of the variation we study and the data on which we rely does not allow us to readily isolate hospital business stealing from these other mechanisms. There were nearly no “untreated” hospitals – all hospitals were affected by the reform and were assigned a revenue shock. The lack of a set of pure controls, in which hospitals were not directly impacted by the reform nor were they impacted indirectly by other hospitals “stealing” their patients, limits our ability to test for business stealing.

A natural, follow-up question is whether these effects on volume and length of stay propagate more broadly through hospitals and affect patients outside the Medicare program. Such “spillovers” can occur when hospitals are not able to treat patients differently based on their insurer, when other insurers mimic the incentives of Medicare, or when hospitals use extra revenue to admit other patients. We investigate this possibility with AHA hospital survey data, which allows us to see patient volume (though not length of stay) by insurer. The next rows of Panel B in Table 2 present estimates on volume based on the AHA data. The third row of Panel B studies Medicare volume in the AHA data. We observe a similar—and, if anything, slightly larger—effect on Medicare admissions when using the AHA data.

The remaining rows of Panel C study other types of admissions: Medicaid admissions and then admissions not covered by either Medicaid or Medicare.⁶ The results suggest that a

⁶The AHA data does not further dis-aggregate patients. Those not covered by Medicaid or Medicare are

positive revenue shock leads to a meaningful increase in *Medicaid* admissions. The magnitude of that effect on Medicaid is nearly the same as the direct effect on Medicare patients, as measured in the AHA data. The point estimate on other admissions is positive though smaller and only statistically significant at the 10-percent level. These findings suggest that the revenue shock induced by Medicare’s payment reform led to broader changes in how hospitals admitted patients, particularly Medicaid patients.

The final panel of Table 2 investigates the mechanisms through which hospitals were able to raise throughput. In particular, we study measures from the AHA survey that relate to each hospital’s ability to scale its operations. We proxy for the size of the hospital’s physical plant and capital inputs by its number of beds. We also consider several measures of labor inputs: full-time equivalent (FTE) employment, nurse employment, and total payroll costs.

The reduced-form estimates suggest that increases in Medicare payments led to statistically significant increases in the number of beds, FTE employment, and nurse employment, and payroll. The 2SLS effects here show that a \$1,000 rise in payment per patient leads to a 3–7-percent rise in the use of these inputs.

Finally, we ask how hospitals’ revenue shocks may have affected the quality of care that they provide. We do so by studying patient outcomes, particularly mortality, given the focus on mortality in clinical trials (see, e.g., Keeley et al., 2003) and value-based payment.

It is challenging in this setting to identify effects on mortality because revenue shocks may have induced a change in the selection of patients appearing in each hospital. Patients are not homogeneous – they differ in the severity of their medical needs. As a result, when a hospital admits more or fewer patients, the characteristics of the marginal patient are likely to change.

This change in the selection of patients admitted to the hospital is both an econometric challenge and a fascinating aspect of hospital behavior in its own right. It is an econometric challenge in that an effect of revenue shocks on mortality could be driven either by hospitals actively changing the quality of care that they provide or simply by a change in the marginal patient. If hospitals increase volume by admitting lower-acuity patients at the margin, then mortality would decrease without a change in quality of care. We attempt to adjust for changes in selection by controlling for patient characteristics, but such an approach only addresses selection on observables. As a result, we must view our analysis of quality as secondary to our focus on volume. That said, we emphasize that either mechanism – a change in selection or a change in quality of care – is interesting and worthy of future study. Indeed, the selection question nests a key health policy issue: when payment changes, *for whom* does access improve or deteriorate?

mostly either privately insured or uninsured.

Appendix Figure A5 describes the effect of the revenue shock on raw and risk-adjusted mortality for three key patient cohorts: heart failure, pneumonia, and acute myocardial infarction (AMI, also known as heart attack).⁷ The figure plots estimates of equation (2) with patient-level data and mortality within 30 days of hospital admission as the outcome. Risk-adjusted estimates control for age, sex, race, and prior health conditions.

The figure provides no suggestion that positive revenue shocks led to an increase in mortality.⁸ Indeed, it suggests a decrease in mortality for patients with heart failure. Evidence on selection is mixed: for AMI, mortality effects attenuate one-fifth with the addition of controls, suggesting that revenue shocks induced hospitals to select observably healthier patients; for the other conditions, we find little sign of selection in either direction.

3 Discussion

A central question in health policy is what and how to reimburse providers. This paper exploits the 2007–2008 Medicare overhaul to isolate how hospitals respond to per-discharge reimbursements. Hospitals that saw large increases in their reimbursement rates reacted by increasing their throughput. The main estimates suggest that for each \$1,000 increase in per-discharge Medicare reimbursement, annual inpatient flows increased by 14 percent.

Another way to describe that finding is by contrasting a static analysis with a dynamic analysis. Label Medicare’s per-visit reimbursement rate as p , and the number of hospitalizations it pays for as q . Policymakers who contemplate a ten-percent increase in Medicare’s per-visit rates would predict that Medicare’s expenditures would change from pq to $1.1 \times pq$. But such an analysis would be a static calculation. The elasticity above suggests that the true effect on Medicare expenditures would be closer to $1.2 \times 1.1 \times pq$.

We can also extrapolate from these results to other changes in hospital reimbursement. CMS’s Hospital-Acquired Condition Reduction Program, for instance, imposes a uniform 1-percent penalty in hospital reimbursement for the bottom-performing quartile of hospitals. The estimates above suggest that a persistently penalized hospital would reduce its Medicare patient volume by roughly 1 percent. To our knowledge, research and debate on the scale of value-based payment has paid little attention to the potential for these programs to change

⁷In the same vein, Appendix Figure A6 describes the effect of the revenue shock on readmissions within 30 days of hospital discharge. The presence of CMS’s concurrent value-based payment program to reduce readmissions, which led to large changes in this outcome (see e.g. Gupta, 2021), makes these estimates more speculative. Hospitals were aware that readmissions would be tracked in this program as of early 2010. The unadjusted effects in Appendix Figure A6 are generally positive, though they attenuate with risk-adjustment.

⁸Mortality effects for patients with pneumonia are less clear; the fungibility of the pneumonia diagnosis has historically led to spurious changes in pneumonia case volume (Silverman and Skinner, 2004), making selection into this cohort a particular concern.

patient volume at penalized or rewarded hospitals.

Finally, further afield, we can extrapolate the estimates to changes in *private* reimbursement. Vogt and Town (2006) suggest that hospital mergers raise commercial prices by 40 percent. This paper's estimates suggest that that increase in prices, in turn, may increase admissions by more than 40 percent and thus more than double total commercial, hospital-based expenditures. Granted, such an extrapolation assumes that effects in Medicare apply to private insurance. But the calculation, if nothing else, raises the possibility that volume effects may matter a great deal whenever the price of a hospitalization changes.

References

- Acemoglu, Daron and Amy Finkelstein**, “Input and technology choices in regulated industries: Evidence from the health care sector,” *Journal of Political Economy*, 2008, 116 (5), 837–880.
- Alexander, Diane and Molly Schnell**, “The Impacts of Physician Payments on Patient Access, Use, and Health,” Technical Report w26095, National Bureau of Economic Research, Cambridge, MA July 2019.
- Azoulay, Pierre, Misty L Heggeness, and Jennifer L Kao**, “Medical Research and Health Care Finance: Evidence from Academic Medical Centers,” Technical Report, National Bureau of Economic Research 2020.
- Baicker, Katherine and Douglas Staiger**, “Fiscal shenanigans, targeted federal health care funds, and patient mortality,” *The quarterly journal of economics*, 2005, 120 (1), 345–386.
- Baker, Laurence C. and Anne Beeson Royalty**, “Medicaid Policy, Physician Behavior, and Health Care for the Low-Income Population,” *The Journal of Human Resources*, 2000, 35 (3), 480.
- Centers for Medicare and Medicaid Services**, “Changes to the Hospital Inpatient Prospective Payment Systems and Fiscal Year 2008 Rates,” *Federal Register*, August 2007, 72 (162), 47129–48175.
- , “Changes to the Hospital Inpatient Prospective Payment Systems and Fiscal Year 2011 Rates,” *Federal Register*, August 2010, 75 (157), 47129–48175.
- , “Design and Development of the Diagnosis Related Group (DRG),” *PBL-038. Baltimore, MD: Centers for Medicare & Medicaid Services*, 2019.
- Chang, Tom and Mireille Jacobson**, “What do nonprofit hospitals maximize? Evidence from California’s seismic retrofit mandate,” 2011. Unpublished.
- 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 Physician Payments,” *Journal of Political Economy*, February 2017, 125 (1), 1–39.
- Cooper, Zack, Amanda Kowalski, Eleanor Powell, and Jennifer Wu**, “Politics and Health Care Spending in the United States,” Technical Report w23748, National Bureau of Economic Research, Cambridge, MA August 2017.
- Currie, Janet and Jonathan Gruber**, “Health insurance eligibility, utilization of medical care, and child health,” *The Quarterly Journal of Economics*, 1996, 111 (2), 431–466.


- **and** – , “Saving babies: The efficacy and cost of recent changes in the Medicaid eligibility of pregnant women,” *Journal of political Economy*, 1996, *104* (6), 1263–1296.
- , – , **and Michael Fischer**, “Physician Payments and Infant Mortality: Evidence from Medicaid Fee Policy,” Technical Report w4930, National Bureau of Economic Research, Cambridge, MA November 1994.
- Cutler, David M**, “The Incidence of Adverse Medical Outcomes Under Prospective Payment,” *Econometrica*, 1995, *63* (1), 29–50.
- **and Richard J Zeckhauser**, “The anatomy of health insurance,” in “Handbook of health economics,” Vol. 1, Elsevier, 2000, pp. 563–643.
- Dafny, Leemore S**, “How do hospitals respond to price changes?,” *American Economic Review*, 2005, *95* (5), 1525–1547.
- Duggan, Mark G**, “Hospital ownership and public medical spending,” *The Quarterly Journal of Economics*, 2000, *115* (4), 1343–1373.
- Frakt, Austin B**, “How much do hospitals cost shift? A review of the evidence,” *The Milbank Quarterly*, 2011, *89* (1), 90–130.
- Frandsen, Brigham, Michael Powell, and James B Rebitzer**, “Sticking points: common-agency problems and contracting in the US healthcare system,” *The RAND Journal of Economics*, 2019, *50* (2), 251–285.
- Gupta, Atul**, “Impacts of performance pay for hospitals: The readmissions reduction program,” 2021.
- Kaestner, Robert and Jose Guardado**, “Medicare reimbursement, nurse staffing, and patient outcomes,” *Journal of Health Economics*, 2008, *27* (2), 339–361.
- Keeley, Ellen C, Judith A Boura, and Cindy L Grines**, “Primary Angioplasty versus Intravenous Thrombolytic Therapy for Acute Myocardial Infarction: A Quantitative Review of 23 Randomised Trials,” *The Lancet*, January 2003, *361* (9351), 13–20.
- MEDPAC**, “Health Care Spending and the Medicare Program,” July 2021.
- Sacarny, Adam**, “Adoption and learning across hospitals: The case of a revenue-generating practice,” *Journal of health economics*, 2018, *60*, 142–164.
- Silverman, Elaine M. and Jonathan S. Skinner**, “Medicare Upcoding and Hospital Ownership,” *J Health Econ*, Mar 2004, *23* (2), 369–89.
- Vogt, William B and Robert Town**, “How has hospital consolidation affected the price and quality of hospital care?,” 2006.
- Wu, Vivian Y and Yu-Chu Shen**, “Long-term impact of Medicare payment reductions on patient outcomes,” *Health services research*, 2014, *49* (5), 1596–1615.

Table 1. Summary Statistics

	Revenue shock > 0		Revenue shock < 0	
Actual DRG weight	1.437	(0.219)	1.400	(0.284)
Avg Medicare payment	9,373	(2,849)	8,012	(2,572)
Avg Medicare charges	32,461	(17,658)	23,602	(13,789)
Avg length of stay, Medicare	5.440	(0.980)	4.783	(0.832)
Beds	273	(194)	215	(189)
Total admissions	13,147	(9,689)	10,094	(9,741)
Medicare admissions	4,340	(3,226)	3,951	(3,430)
Medicaid admissions	2,484	(2,497)	1,636	(1,988)
Fraction Medicare	0.424	(0.102)	0.474	(0.095)
Total annual payroll	82	(87)	61	(79)
Full time RNs	415	(390)	312	(361)
Full time equivalents	1,512	(1,419)	1,187	(1,342)
Salary (total payroll/FTEs)	53,154	(12,853)	48,293	(12,378)
Share Non profit	0.69		0.67	
Share For profit	0.17		0.17	
Share Government	0.14		0.16	
Share Urban	0.87		0.62	
Hospitals	1,112		1,396	

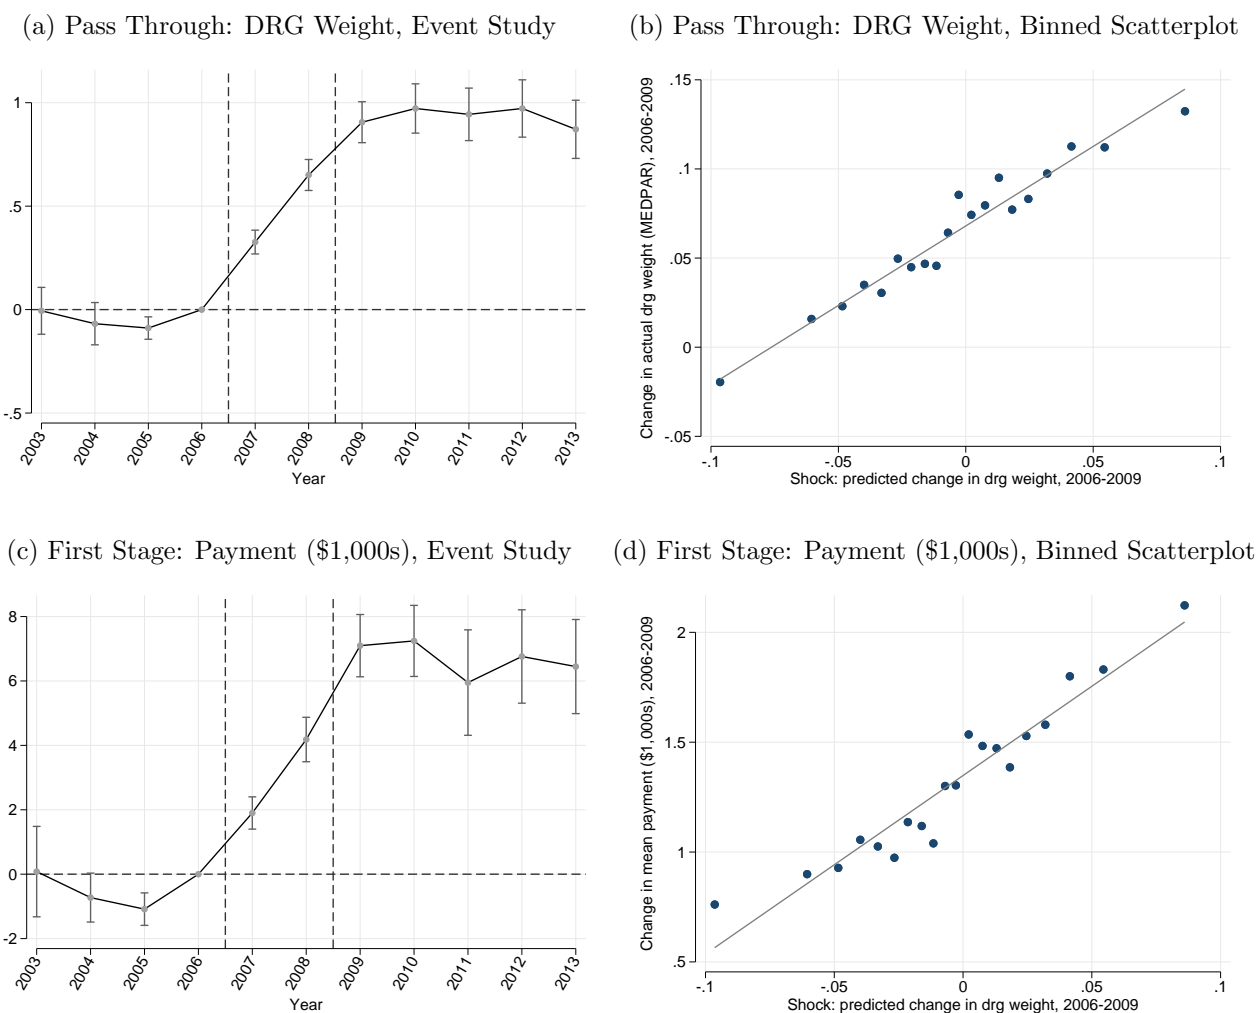
This table presents 2006 summary statistics dividing hospitals by the shock instrument. Means are presented with standard deviations in parentheses. Hospitals with a shock below zero were set to experience lower reimbursement rates. Hospitals with a shock above zero were set to experience higher reimbursement rates. Total annual payroll measured in millions.

Figure 1. Determinants and Distribution of the Revenue Shock



These figures show the linkage between a hospital’s share of patients with major complications and comorbidities (MCCs) and its revenue shock. They plot (a) a histogram of the hospital-specific share of patients with an MCC in 2006, (b) a histogram of the hospital-specific revenue shock, calculated as the difference between a hospital’s DRG weight from running 2009 MS-DRG grouper on the 2006 claims and its actual 2006 average DRG weight (equation 1), and (c) a binned scatterplot of the relationship between each hospital’s 2006 MCC share and its revenue shock.

Figure 2. Pass Through and First Stage



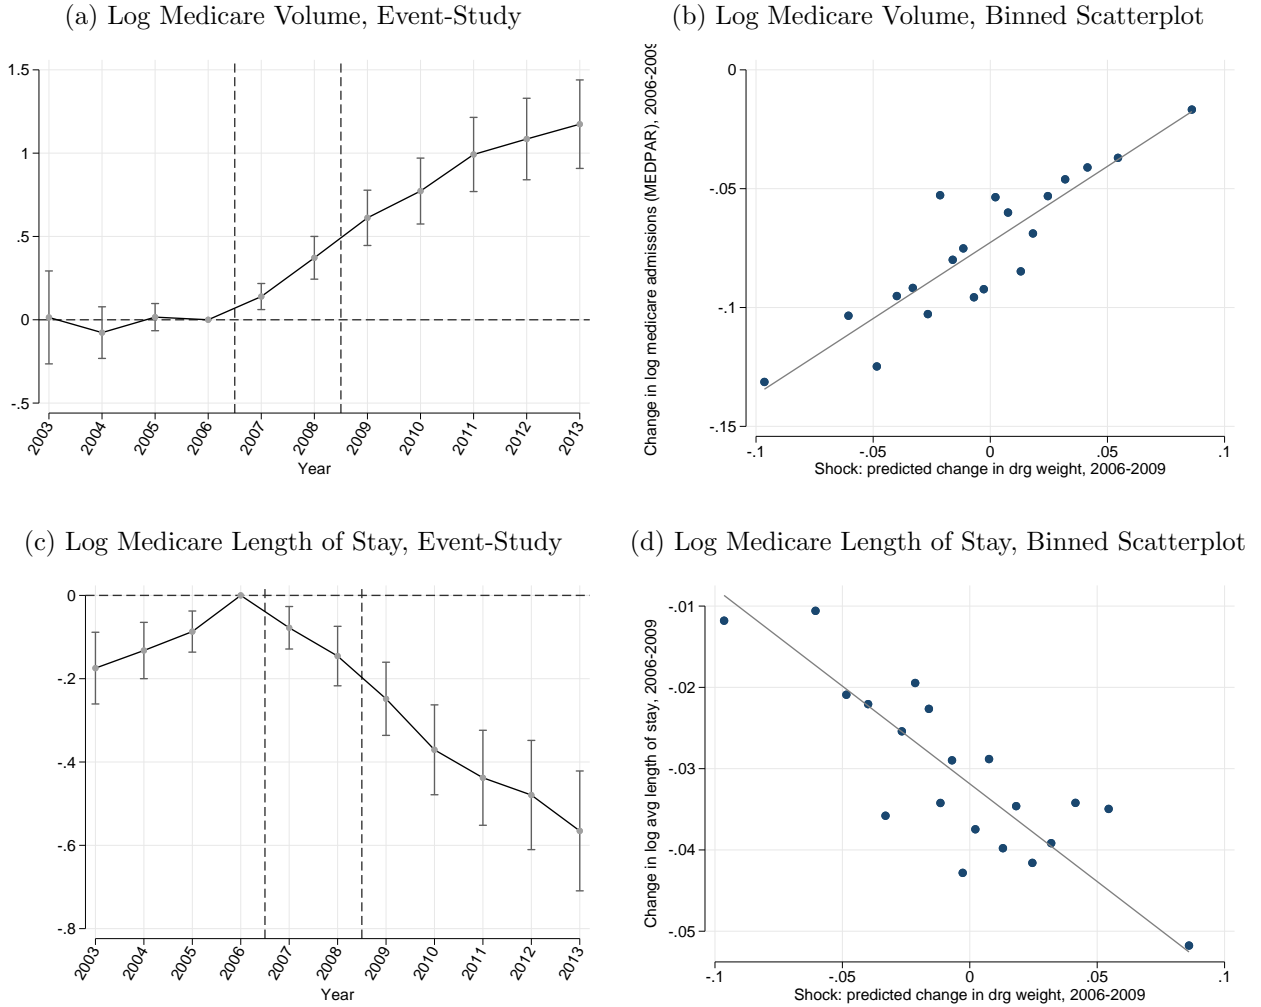
These figures describe the relationship between the revenue shock and the actual DRG weight and Medicare payments received by hospitals. The first and third panels present estimates of the β_s coefficients from equation (2) for each year. Hospital-specific fixed effects, year-specific fixed effects, and hospital-size-times-year-specific fixed effects not shown. The 95-percent confidence intervals are based on standard errors clustered at the level of the hospital. The omitted year is 2006. The second and fourth panels present binned scatterplots with the actual change in hospital reimbursements plotted along the vertical axis, and the revenue shock on the horizontal axis. The sample is restricted to hospitals with at least 500 Medicare patients a year in 2004–2013 and excludes critical-access hospitals.

Table 2. Impact of Revenue Shock on First Stage, Hospital Volume, and Resources – un-weighted

	Reduced Form	IV	<i>N</i>	Hospitals
A. Pass Through				
DRG Weight	0.974*** (0.061)		27,575	2,508
B. First Stage				
Medicare Payment (\$1,000s)	7.134*** (0.667)		27,575	2,508
C. Logged Hospital Volume and Throughput				
Medicare Volume (MEDPAR data)	0.939*** (0.125)	0.130*** (0.020)	27,575	2,508
Medicare Length of Stay	-0.322*** (0.061)	-0.045*** (0.010)	27,575	2,508
Medicare Volume (AHA data)	1.202*** (0.135)	0.170*** (0.024)	27,263	2,502
Medicaid Volume	0.794*** (0.223)	0.112*** (0.033)	27,236	2,502
Other Volume	0.308* (0.179)	0.043* (0.025)	27,282	2,504
D. Logged Hospital Scale Measures				
Beds	0.330*** (0.124)	0.046*** (0.018)	27,326	2,504
Full-time Equivalentents	0.357*** (0.117)	0.050*** (0.016)	27,325	2,504
Nurses Employed	0.497*** (0.125)	0.070*** (0.018)	27,310	2,504
Payroll	0.239** (0.116)	0.034** (0.016)	27,326	2,504

In panels C and D, the logarithm of the given variable is the outcome of interest. The first and second panels present estimates of equation (3) with the listed variables as the outcomes of interest. Hospital-specific fixed effects, year-specific fixed effects, shock-times-interim effect, and hospital-size-times-year interaction terms not shown. The third panel and fourth panels present 2SLS estimates in which we instrument for average Medicare payments, measured in thousands of dollars, with the revenue shock times post. Hospital-specific fixed effects, year-specific fixed effects, shock-times-interim effect, and hospital-size-times-year interaction terms not shown. Standard errors in parentheses clustered on hospital. The sample is restricted to hospitals with at least 500 Medicare patients a year in 2004–2013 and excludes critical-access hospitals. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 3. Reduced-Form Effects



These figures describe the relationship between the revenue shock and Medicare admissions. The first and third panels present estimates of the β_s coefficients from equation (2) for each year with the variables listed as the outcomes of interest. Hospital-specific fixed effects, year-specific fixed effects, and hospital-size-times-year interaction terms not shown. The 95-percent confidence intervals are based on standard errors clustered at the level of the hospital. The omitted year is 2006. The second and fourth panels present a binned scatterplot with the change in log Medicare admissions or log Medicare length of stay plotted along the vertical axis, and the revenue shock plotted on the horizontal axis. The sample is restricted to hospitals with at least 500 Medicare patients a year in 2004–2013 and excludes critical-access hospitals.

Appendix To:
Regulated Revenues and Hospital Behavior: Evidence
from a Medicare Overhaul

Tal Gross, Adam Sacarny, Maggie Shi, and David Silver

August 2022

For Online Publication

A Appendix: Review of Hospital Pay-For-Performance Programs

This section provides an overview of recent hospital payment reforms. As mentioned in the main text, the Patient Protection and Affordable Care Act (ACA) made a handful of key payment reforms to align hospital payment and quality of care: the Hospital Readmissions Reduction Program (HRRP), the Hospital Value-Based Purchasing Programs (HVBPP), and the Hospital-Acquired Conditions Reduction Program (HACRP).

There are a few features of these reforms that are important for understanding their impacts on hospitals' payment rates. First, each of these programs administers payment adjustments as percent modifiers of Medicare reimbursements for a given fiscal year. They do so through what is known as a hospital's base operating payment. For instance, under HRRP in fiscal years 2017 and later, each hospital faces up to a 3-percent cut in total Medicare reimbursements through an adjustment of its base operating payment, despite the fact that currently only 6 conditions (3 conditions in FY2013) are considered in determining annual hospital-specific penalties.

Second, payment adjustments in each program are determined based on rank orderings of hospitals over a given period. Combined with imperfect risk adjustment that disadvantages hospitals that treat patients with unmeasured and/or uncoded risks, rank-based penalties have led to criticisms regarding fairness and the ability of hospitals to improve enough in their rankings to avoid penalty. CMS has made some adjustments to address this criticism. For example, in 2019 CMS introduced peer-comparison groups under HRRP, so that hospitals are compared to peers with similar shares of Medicare-Medicaid dual eligible patients, a proxy for socioeconomic risk factors. Despite that change, the equity of these programs across hospitals remains a concern.

Third, CMS uses a multi-year (often 2- or 3-year) lookback period to measure performance in each domain. The multi-year lookback period is meant to improve the reliability of program metrics such as readmission rates for specific, sometimes rare conditions. Thus, payment adjustments from one year to the next rely on overlapping lookback periods, creating a mechanical "stickiness" of payment adjustments over time.

As a result of these program features, hospital payment adjustments stemming from the ACA's performance-pay programs were predictable and persistent. Hospitals with difficult patient mixes bore larger penalties both at program inception and today. Totaled across the three programs, payment adjustments were correlated approximately 0.8 year-to-year, and 0.3 between FY2013 and FY2018.


These pay-for-performance reforms had a much smaller impact on hospital reimbursement

rates than the transition to MS-DRGs that we study. First the reforms were smaller in variability across hospitals. The standard deviation of total adjustments from these programs was 0.5 percent in FY2013 and reached 1.1 percent in FY2018. By comparison, the 2008 switch to MS-DRGs induced payment rate adjustments with a standard deviation of 5.2 percent across hospitals.

Second, the reforms were smaller in terms of the total amount of money shifted across hospitals. CMS estimated that the transition to MS-DRGs would re-distribute 2 percent of inpatient prospective payments, which would amount to about \$2 billion in reimbursements redistributed across hospitals. By contrast, total HRRP penalties were well under half a billion dollars in 2015.

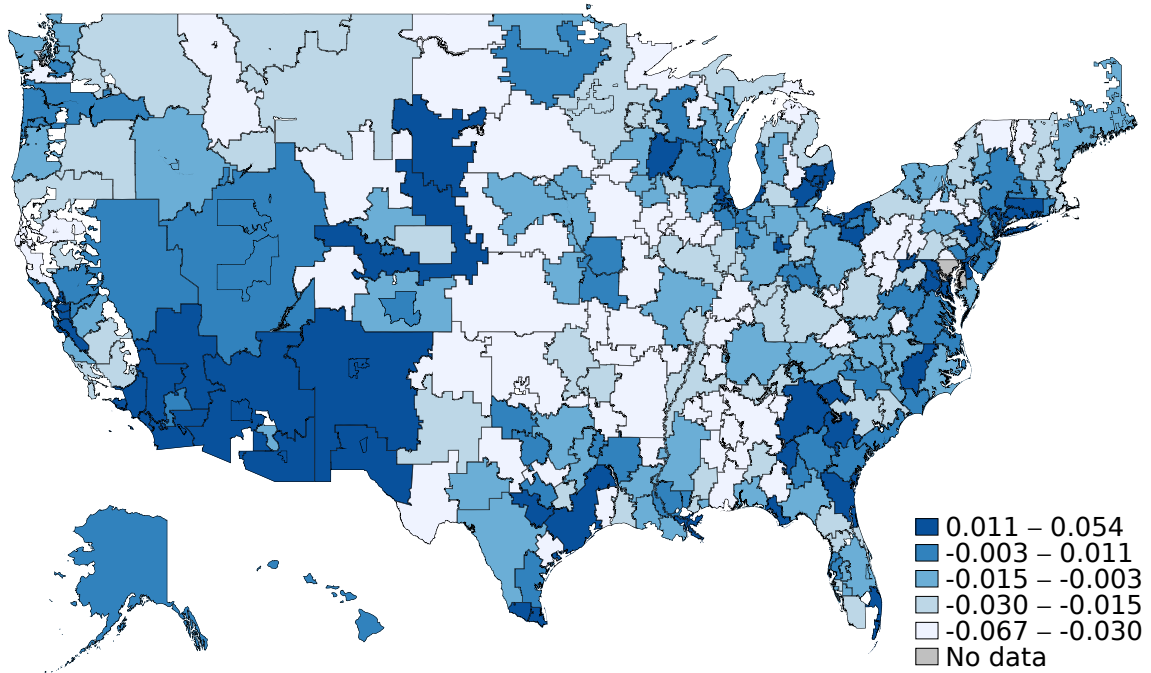
B Appendix Figures

Figure A1. Average DRG Weights for Hospitals That Experienced The Largest Changes in Reimbursement




This figure plots the average of actual relative prices received by hospitals from Medicare. We restrict the sample to two groups of hospitals: those in the top decile of the revenue shock and those in the bottom decile. The sample is restricted to hospitals with at least 500 Medicare patients a year in 2004–2013 and excludes critical-access hospitals.

Figure A2. Geographic Distribution of Revenue Shock by HRR



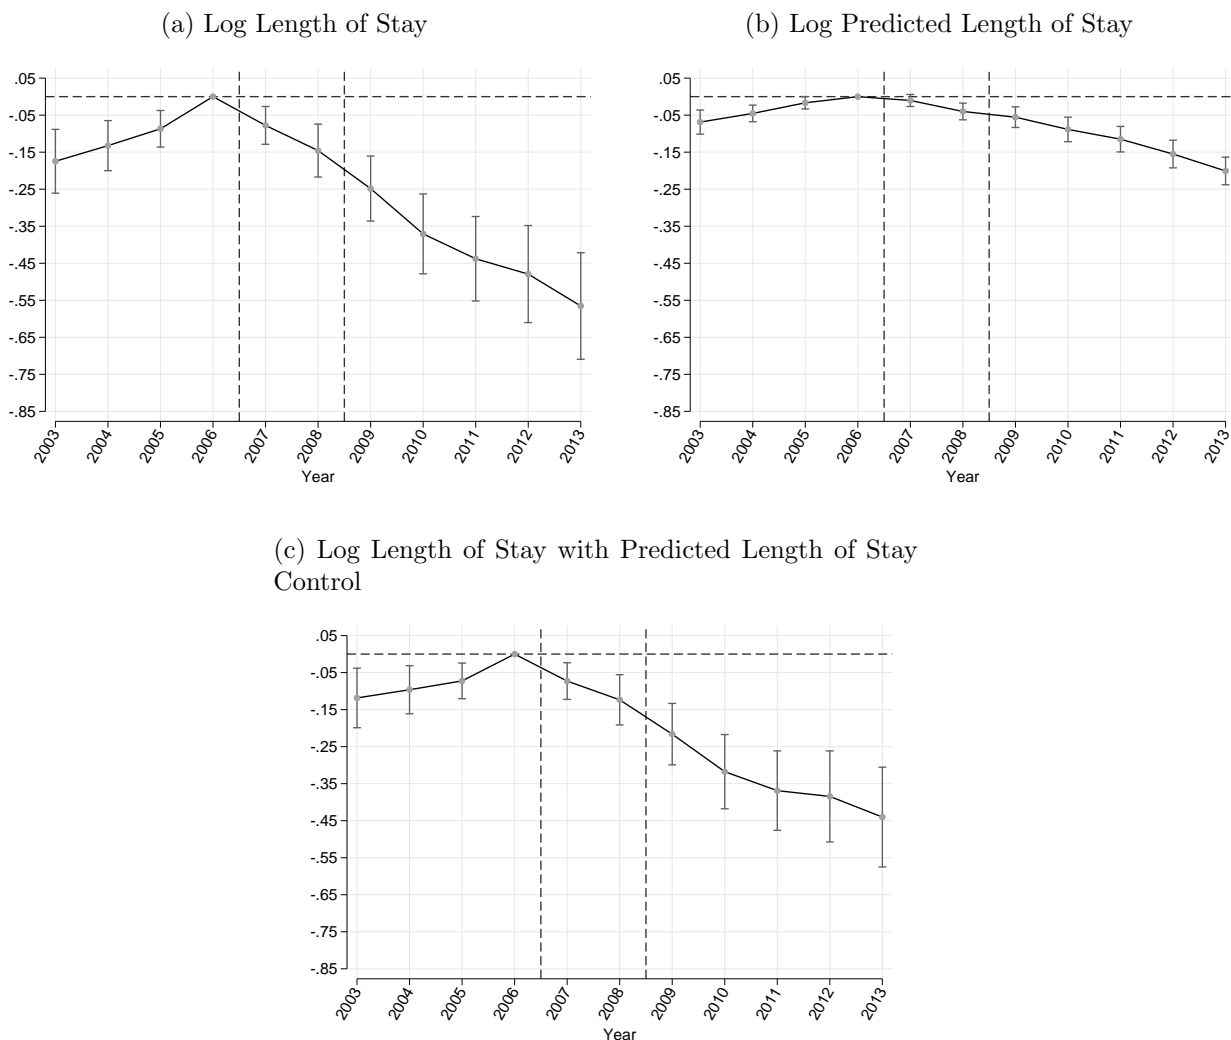
This figure plots the average of our key measure of the change in Medicare reimbursements by hospital referral region (HRR). The sample consists of 2,498 hospitals and is restricted to hospitals with at least 500 Medicare patients a year in 2004–2013 and excludes Critical Access Hospitals. There are 304 HRRs in total.

Figure A3. Time Series of Stroke Reimbursement Rate



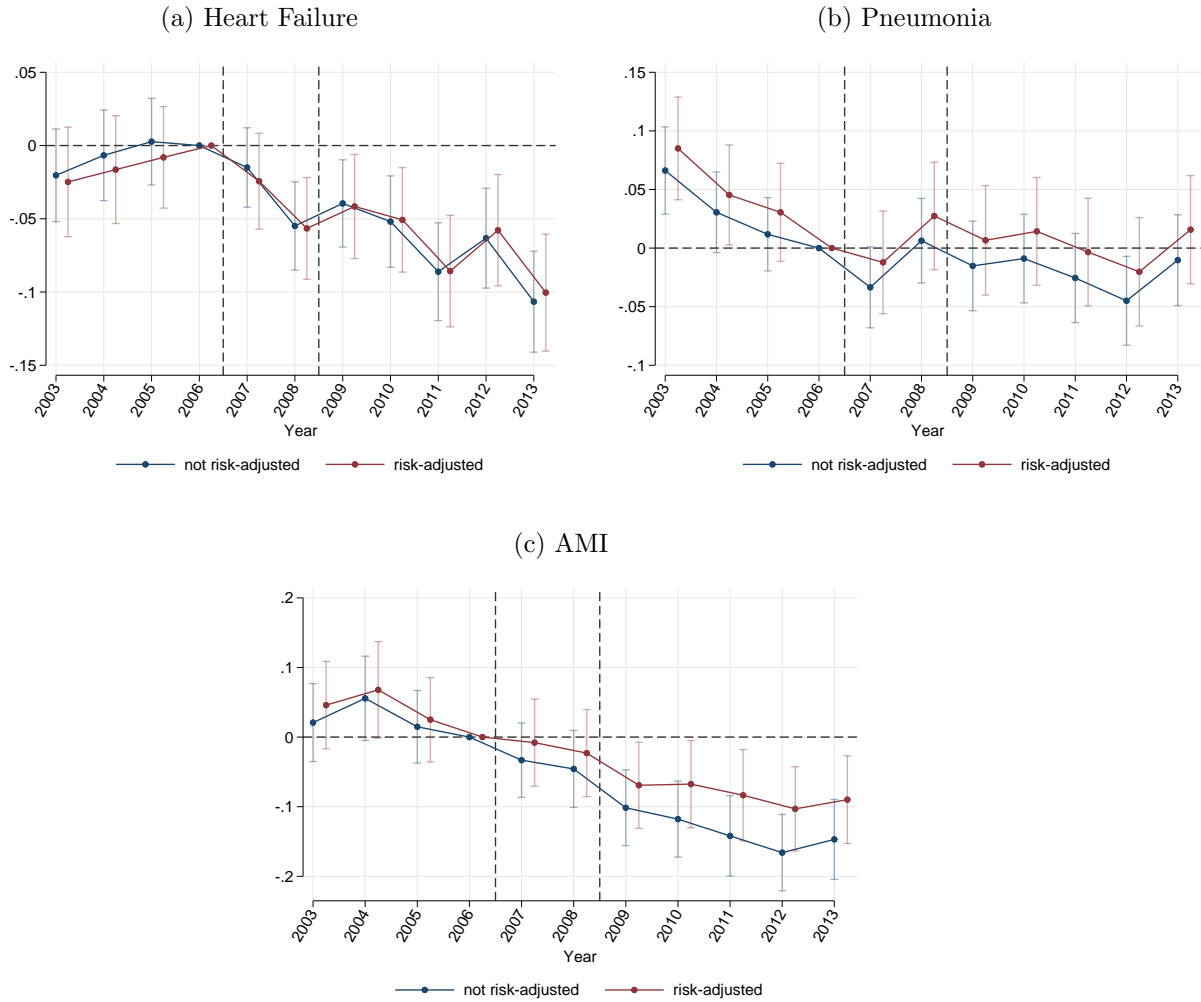
This figure plots the relative price paid by Medicare to hospitals across years for one particular, common condition: stroke (DRG 14 in 2007, labeled intracranial hemorrhage or cerebral infarction). The figure illustrates the policy change, which shifted from one level of severity through FY2007 to three levels of severity after 2008.

Figure A4. Reduced Form: Length of Stay and Predicted Length of Stay




Figures (a) and (b) present estimates of the β_s coefficients from equation (2) when the outcome of interest is Medicare length of stay and predicted length of stay. Hospital-specific fixed effects and year-specific fixed effects not shown. Length of stay is predicted using a patient's race, age, sex, chronic conditions, and major diagnostic category using pre-reform data. Figure (c) presents estimates of the coefficients from including predicted length of stay as a control in equation (2), where the outcome of interest is length of stay. The 95-percent confidence intervals are based on standard errors clustered at the level of the hospital. The sample is restricted to hospitals with at least 500 Medicare patients a year in 2004–2013 and excludes critical-access hospitals. The omitted year is 2006.

Figure A5. Effect of the Revenue Shock on Patients' Mortality



These figures present estimates of the β_s coefficients from equation (2) when 30-day mortality of index admissions for the given condition is the outcome of interest. Hospital-specific fixed effects and year-specific fixed effects not shown. The risk-adjusted estimates control for age, sex, race, and prior health conditions. The 95-percent confidence intervals are based on standard errors clustered at the level of the hospital. The sample is restricted to hospitals with at least 500 Medicare patients a year in 2004–2013 and excludes critical-access hospitals. The omitted year is 2006.

Figure A6. Effect of the Revenue Shock on Re-admissions



These figures present estimates of the β_s coefficients from equation (2) when 30-day revisits of index admissions for the given condition is the outcome of interest. Hospital-specific fixed effects and year-specific fixed effects not shown. The risk-adjusted estimates control for age, sex, race, and prior health conditions. The 95-percent confidence intervals are based on standard errors clustered at the level of the hospital. The sample is restricted to hospitals with at least 500 Medicare patients a year in 2004–2013 and excludes critical-access hospitals. The omitted year is 2006.

C Appendix Tables

Table A1. Effects of the Revenue Shock by Hospital Ownership

	(1)	(2)	(3)	(4)	(5)
	Average Total Payment (\$1000s)	Log Medicare Volume		Log Medicare Length of Stay	
For profit \times Revenue Shock \times Post	8.356*** (1.623)	0.816*** (0.281)		-0.781*** (0.150)	
Non profit \times Revenue Shock \times Post	6.609*** (0.825)	0.860*** (0.147)		-0.193*** (0.063)	
Government \times Revenue shock \times Post	6.770*** (1.875)	1.356*** (0.329)		-0.183 (0.137)	
For profit \times Average Total Payment (\$1000s)			0.0931*** (0.0358)		-0.0935*** (0.0254)
Non profit \times Average Total Payment (\$1000s)			0.130*** (0.0266)		-0.0292** (0.0114)
Government \times Average Total Payment (\$1000s)			0.200*** (0.0703)		-0.0268 (0.0221)
For-profit v. Non-profit	0.34	0.89	0.41	0.00**	0.02**
For-profit v. Government	0.52	0.21	0.18	0.00**	0.05**
Non-profit v. Government	0.94	0.17	0.35	0.94	0.92
Hospitals	2,508	2,508	2,508	2,508	2,508
N	27,575	27,576	27,575	27,575	27,575

Total payment is measured in thousands of dollars and the outcomes for columns 2 through 5 are the logarithm of Medicare admissions and the logarithm of the average length of stay. Columns 1, 2, and 4 present estimates of equation (3). Columns 3 and 5 present 2SLS estimates, in which we instrument for total payment with the revenue-shock instrument interacted with a post-2006 indicator function. The table lists p -values that test equality across interaction terms: for-profit versus non-profit; for-profit versus government; and non-profit versus government. Hospital-specific fixed effects, year-specific fixed effects, shock-times-interim effect, and hospital-size-times-year interaction terms not shown. The sample is restricted to hospitals with at least 500 Medicare patients a year in 2004–2013 and excludes critical-access hospitals. Standard errors in parentheses clustered on hospital. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A2. Robustness, Length of Stay

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. First-Stage Outcome: Average Total Payment							
Revenue Shock × Post	7.134*** (0.667)	6.574*** (0.724)	6.823*** (0.771)	7.780*** (0.656)	4.267*** (0.759)	4.204*** (0.858)	6.848*** (0.659)
B. Reduced-Form Outcome: Log Medicare Length of Stay							
Revenue Shock × Post	-0.322*** (0.0613)	-0.412*** (0.0725)	-0.400*** (0.0711)	-0.323*** (0.0613)	-0.344*** (0.0776)	-0.357*** (0.0843)	-0.315*** (0.0593)
C. 2SLS Outcome: Log Medicare Length of Stay							
Average Total Payment	-0.0452*** (0.0104)	-0.0621*** (0.0146)	-0.0591*** (0.0142)	-0.0415*** (0.00933)	-0.0806*** (0.0249)	-0.0850*** (0.0290)	-0.0460*** (0.0104)
Shock year	2006	2005	2004	2006	2006	2006	2006
Bed bin × year	✓	✓	✓		✓	✓	✓
HRR × year					✓		
CZ × year						✓	
County MA pen.							✓
Hospitals	2,508	2,508	2,508	2,508	2,496	2,333	2,504
<i>N</i>	27,575	27,575	27,575	27,575	27,443	25,433	27,305

This table presents estimates of equation (3) with the logarithm of average length of stay as the outcomes of interest. Hospital-specific fixed effects, year-specific fixed effects, and shock-times-interim terms not shown. Standard errors in parentheses clustered on hospital. The sample is restricted to hospitals with at least 500 Medicare patients a year in 2004–2013 and excludes critical-access hospitals. All payments are measured in thousands of dollars. “CZ” indicates commuting zones, “HRR” indicates hospital referral region; “MA pen.” indicates yearly Medicare Advantage penetration.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A3. Robustness, Volume

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. First-Stage Outcome: Average Total Payment							
Revenue Shock × Post	7.134*** (0.667)	6.574*** (0.724)	6.823*** (0.771)	7.781*** (0.656)	4.267*** (0.759)	4.204*** (0.858)	6.849*** (0.659)
B. Reduced-Form Outcome: Log Medicare Volume							
Revenue Shock × Post	0.939*** (0.125)	0.862*** (0.133)	0.569*** (0.127)	1.075*** (0.117)	0.503*** (0.134)	0.366*** (0.139)	0.855*** (0.119)
C. 2SLS Outcome: Log Medicare Volume							
Average Total Payment	0.130*** (0.0203)	0.125*** (0.0231)	0.0811*** (0.0195)	0.137*** (0.0179)	0.118*** (0.0374)	0.0871** (0.0369)	0.125*** (0.0206)
Shock year	2006	2005	2004	2006	2006	2006	2006
Bed bin × year	✓	✓	✓		✓	✓	✓
HRR × year					✓		
CZ × year						✓	
County MA pen.							✓
Hospitals	2,508	2,508	2,508	2,508	2,496	2,333	2,504
<i>N</i>	27,576	27,576	27,576	27,576	27,444	25,433	27,305

This table presents estimates of equation (3) with the logarithm of Medicare admissions as the outcomes of interest. Hospital-specific fixed effects, year-specific fixed effects, and shock-times-interim terms not shown. Standard errors in parentheses clustered on hospital. All payments are measured in thousands of dollars. “CZ” indicates commuting zones, “HRR” indicates hospital referral region; “MA pen” indicates yearly Medicare Advantage penetration. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$