How Effective is Ex Ante Enforcement? Evidence From 15 Billion Medicare Claims*

John Griffin † Christoph Herpfer‡ Gonzalo Maturana§

December 11, 2025

[PRELIMINARY DRAFT]

Ex ante administrative enforcement is a widely used method to prevent fraud, but there is little evaluation of its effectiveness. We study the effectiveness of one of Medicare's main administrative programs for hospital coding: the staggered rollout of Medically Unlikely Edits (MUEs), which set objective upper bounds on how often each procedure may be billed. Before these limits take effect, hospitals submit claims that exceed future thresholds at a rate of 1.2% of total claims. This behavior is highly concentrated, with the top 10% of hospitals accounting for 87% of all violations and exhibiting violation rates several times higher than typical providers. When MUEs become enforceable, violations on treated codes fall by about 95%, producing visible reductions in upcoding. However, violations quickly migrate to codes without limits and later to "day-of-service" violations created by splitting charges across claim lines, and day-of-service limits. Overall, aggregate noncompliant billing declines only modestly because hospitals systematically adjust their billing to evade administrative controls, with total annual damage to CMS exceeding $1 billion in 2023 despite all enforcement measures. This highlights both the prevalence of fraud in healthcare and the challenges of curbing it through rule-based enforcement by government regulators.

JEL classification: H51, I13, I18

Keywords: Medicare, medical billing, medically unlikely edits (MUE), healthcare regulation

*We thank seminar participants at the University of Iowa for their helpful comments. Integra Med Analytics provided data assistance and research support. The firm specializes in medical fraud detection and is involved in litigation related to fraudulent practices in the healthcare sector, including skilled nursing and hospitals. Integra Med Analytics also provides consulting services on healthcare issues. Griffin is the owner and CEO of Integra Med Analytics, Integra FEC, and Integra Research Group.

†Texas McCombs School of Business. Email: john.griffin@utexas.edu.

‡UVA Darden School of Business. Email: herpferc@darden.virginia.edu (corresponding author).

§Emory Goizueta Business School. Email: gonzalo.maturana@emory.edu.

How Effective is Ex Ante Enforcement? Evidence From 15 Billion Medicare Claims

December 11, 2025

[PRELIMINARY DRAFT]

Ex ante administrative enforcement is a widely used method to prevent fraud, but there is little evaluation of its effectiveness. We study the effectiveness of one of Medicare's main administrative programs for hospital coding: the staggered rollout of Medically Unlikely Edits (MUEs), which set objective upper bounds on how often each procedure may be billed. Before these limits take effect, hospitals submit claims that exceed future thresholds at a rate of 1.2% of total claims. This behavior is highly concentrated, with the top 10% of hospitals accounting for 87% of all violations and exhibiting violation rates several times higher than typical providers. When MUEs become enforceable, violations on treated codes fall by about 95%, producing visible reductions in upcoding. However, violations quickly migrate to codes without limits and later to "day-of-service" violations created by splitting charges across claim lines, and day-of-service limits. Overall, aggregate noncompliant billing declines only modestly because hospitals systematically adjust their billing to evade administrative controls, with total annual damage to CMS exceeding $1 billion in 2023 despite all enforcement measures. This highlights both the prevalence of fraud in healthcare and the challenges of curbing it through rule-based enforcement by government regulators.

JEL classification: H51, I13, I18

Keywords: Medicare, medical billing, medically unlikely edits (MUE), healthcare regulation

There is increasing evidence that fraud in government spending, and in Medicare in particular, is widespread,[1] yet there is little consensus on how to contain it. Enforcement relies on two main approaches: ex post legal actions pursued through civil or criminal cases, and ex ante administrative controls that attempt to block improper payments in real time. Government programs during the pandemic are thought to have mostly utilize the former method of enforcement, which is often colloquially described as "pay now, chase later" and allowed approximately $300 and $429 billion in pandemic fraud.[2] Medicare has expanded its use of administrative tools, but there is limited evidence on whether they are effective. One of the largest such programs is the set of daily, per-patient ex ante billing limits known as Medically Unlikely Edits (MUEs), established by the Centers for Medicare & Medicaid Services (CMS) in collaboration with the American Medical Association (AMA), which place procedure-level caps on how often a service can be billed. In this paper, we provide the first comprehensive evaluation of this program and examine how providers respond to these limits.

The healthcare sector accounts for roughly 20% of U.S. GDP and Medicare is its largest payer, providing coverage primarily for older adults (American Medical Association, 2024). The structure of this market makes it especially vulnerable to improper billing: patients rarely pay directly and have limited ability to monitor care, while insurers cannot observe the patient-provider interaction and therefore cannot verify what was actually delivered. These features create classic agency problems (Arrow, 1963; Jensen and Meckling, 1976) and open the door to subtle forms of overbilling that are difficult to identify. Although statistical tools can highlight providers with implausible billing patterns (Li, 2008; Centers for Medicare & Medicaid Services, 2024), they also underscore the need for settings that

[1]Annual losses to fraud in healthcare are estimated to range from $100 billion to more than $300 billion (National Health Care Anti-Fraud Association 2025; U.S. Department of Justice 2025).

[2]Griffin, Kruger, and Mahajan (2025) detail how this type of grassroots fraud spread rapidly through social networks to outpace government enforcement. They collect pandemic estimates of fraud from government reports and academic research that range between $300 and $429 billion or 14.5% to 20.6% of the $2.07 trillion spent in the three government programs (PPP, EIDL, and unemployment insurance).

provide clear, verifiable indicators of improper billing.

Our approach overcomes the usual challenges in detecting fraudulent Medicare billing by exploiting the fact that MUEs create objective, medically determined daily limits on how frequently each procedure can be billed. These limits provide a clear benchmark for identifying implausible billing. For intuition, certain procedures such as overnight sleep studies cannot be performed more than once per day, so billing multiple units for the same patient is inherently implausible. We combine the universe of Medicare Part B claim-line records from 2012 to 2023 with a hand-constructed panel tracking the introduction of new MUE thresholds. This setting allows us to measure how much excessive billing occurred before limits were published and to study whether providers reduced or adapted their billing once limits were in place. Because violations represent clear departures from medically plausible practice, our estimates provide a conservative lower bound on improper billing that exceeds $1 billion per year for this specific form of overbilling. The data also allow us to identify the providers involved and to document how billing behavior changes in response to enforcement

Our empirical setup exploits the staggered introduction of MUE limits across billing codes. Because we observe the exact date on which each limit becomes effective, we can examine billing behavior for procedures that were later assigned an MUE. This allows us to identify what we refer to as *pre-MUE violations*, meaning instances in which providers billed units that exceed the future MUE threshold before the limit was officially published. These violations are defined retrospectively once the thresholds become known. Although they were not automatically flagged at the time, they represent billing patterns that would later be deemed medically implausible. We then track whether providers reduce, persist in, or adapt their billing behavior once the limits become enforceable, focusing on the same procedure codes before and after the introduction. Finally, we use a difference-in-differences framework to account for unobserved hospital-level characteristics that may influence changes in billing over time.

Pre-MUE violations are associated with substantial costs to the government, totaling approximately $15 billion in improper payments during our sample period from 2012 to 2024, with average quarterly payments of about $300 million. These violations are highly unevenly distributed across hospitals. A large share of providers never exceed these limits, yet a small subset consistently bills far above what is medically plausible. The top 1% of hospitals account for 33% of all pre-MUE violations, and the top 10% account for 87%. This concentration suggests that a limited group of providers specialize in aggressive billing practices that generate a disproportionate share of improper charges. MUE violations also display strong persistence. Hospitals with a high rate of pre-MUE violations in one quarter tend to remain high violators in subsequent quarters, suggesting that this behavior reflects systematic practice rather than temporary opportunism.

To test whether hospitals exceed limits because they treat patients with unusually high medical needs, we examine how their billing behavior changes once MUE limits become enforceable. In principle, some violations could be clinically justified if hospitals specialize in patients who legitimately require more intensive or frequent procedures. However, the evidence is inconsistent with this explanation. Using a difference-in-differences framework that compares the same procedure before and after its limit is introduced while accounting for hospital-level shocks, we find that violations fall sharply once the MUE threshold becomes binding. The fraction of claim lines that exceed the limit drops from more than 1% to less than 0.25% within two years, a decline of roughly 75%. Since hospitals can still bill above the limit when they provide a valid medical justification, this abrupt and persistent reduction indicates that pre-MUE violations are not driven by patient complexity but reflects intentional and profit-oriented behavior.

Next, we consider a benign alternative explanation for the drop in violations: that the observed reduction reflects honest accounting errors rather than intentional evasion. We evaluate this possibility using three complementary tests.

The first test examines substitution. If violations were primarily unintentional, hos-

pitals should reduce questionable billing broadly once an MUE threshold becomes salient for a specific code. In contrast, intentional behavior implies shifting activity toward codes that remain unconstrained. The evidence strongly supports the latter interpretation. Over the sample period, the share of procedural codes with an MUE assignment increases from roughly 70% to 99%, leaving only a small and shrinking set of uncovered codes where violations remain mechanically possible. As this unconstrained set contracts, the violation rate within it rises sharply, from roughly 1–2% early in the sample to more than 10% in 2017, before stabilizing around 5%. By the end of the period, the remaining uncovered codes largely correspond to newly introduced procedures not yet assigned an MUE. This pattern is difficult to reconcile with innocent error and is instead consistent with purposeful substitution in response to enforcement.

The second test examines persistence. If violations stem from honest mistakes, hospitals' violation rates should vary substantially over time. Instead, violation behavior is highly persistent. A formal test shows that the AR(1) coefficient of the fraction of violating claims is large and statistically significant at 97.5%. Such stability is consistent with deliberate, systematic billing practices rather than occasional lapses or random errors.

The third test considers financial incentives. If CMS routinely denied claims that exceeded MUE limits, violations might simply reflect clerical errors with no economic consequences. Instead, violating claims are regularly paid before MUE limits are introduced, and payments decline only once the limits become enforceable. In total, these violations account for roughly $1 billion per year in Medicare payments. This pattern is again inconsistent with accidental errors and aligns with the view that hospitals respond to economic incentives.

Taken together, these three tests provide strong evidence that the observed behavior reflects intentional strategies to sustain excessive billing despite regulatory constraints, rather than inadvertent mistakes.

We next provide additional evidence consistent with intentional behavior. Violations are

disproportionately concentrated in more lucrative procedures: although they represent only 1.2% of submitted claims, they account for 1.5% of total charges. A regression that relates violations to average reimbursement levels confirms this pattern. Even when comparing the same procedure code over time, higher reimbursement is associated with significantly higher violation rates.

A final and particularly compelling piece of evidence comes from hospitals' responses to the introduction of day-of-service (DoS) limits. Once claim-line MUEs make certain violations mechanically impossible, hospitals do not reduce questionable billing broadly. Instead, they switch margins: rather than billing excessive units on a single claim line, they split the same procedure across multiple lines within the day, avoiding the claim-line edit while still exceeding the medically plausible daily total. Because DoS limits cap the total number of units that can be billed per patient per day, this behavior constitutes a new violation rather than clinical necessity.

We assemble a panel of DoS limits and track these violations over time. DoS violations rise sharply when claim-line violations fall, offsetting roughly 95% of the reduction achieved by the claim-line edits. Moreover, DoS limits themselves have little detectable deterrent effect: a difference-in-differences analysis shows no meaningful decline in DoS violations charged after their introduction, pointing to weak enforcement and substantial room for evasion.

Taken together, these results show that hospitals actively re-optimize their billing strategies when constraints tighten, indicating purposeful and economically motivated behavior rather than clerical errors. They also highlight the inherent limits of administrative regulation in a complex system with multiple avenues for circumvention.

Our paper contributes to the literature on the effectiveness of anti-fraud enforcement. A small but growing set of studies evaluates how Medicare's tools deter improper billing. Eliason et al. (2021) show that requiring prior authorization for ambulance transports to dialysis centers sharply reduced unnecessary rides and was more effective than criminal

or civil enforcement. Our findings qualify this result by showing that administrative prohibitions can be effective in narrow settings but are quickly circumvented when providers can switch to alternative codes. Leder-Luis (2025) document that False Claims Act actions generate deterrence effects roughly ten times larger than the settlement amounts, consistent with related evidence in Howard and McCarthy (2020) and Flasher and Lamboy-Ruiz (2019). Shi (2023) find that Medicare's Recovery Audit Contractor (RAC) program, although applied to only 4% of claims, produces future savings 24 to 29 times its cost, largely through deterrence. Our results complement this work by demonstrating that MUE limits, a widely used administrative tool, have weak deterrent effects for serial violators and are easily gamed, which suggests that more intensive RAC auditing and Department of Justice (DOJ) enforcement may be necessary to curb persistent abuse.

Second, our paper relates to the broad literature on healthcare fraud, including upcoding and medically unnecessary care. Leder-Luis and Malani (2025) review these practices and classify fraud into upcoding and the provision of unnecessary services. Hospital upcoding has been documented in many settings (Silverman and Skinner, 2004; Dafny, 2005; Heese et al., 2015; Bastani et al., 2019; Joiner et al., 2024; Shekhar et al., 2023). Geruso and Layton (2020) show that Medicare Advantage plans report risk scores 6–16% higher than fee-for-service Medicare for identical patients, implying substantial excess spending. Griffin and Priest (2024) document large differences in upcoding in skilled nursing under different reimbursement regimes.[3] Related work studies broader billing practices and organizational influences on pricing (Howell et al., 2021; Liu, 2022; Eliason et al., 2020; Andreyeva et al., 2023). Evidence of unnecessary treatment has been shown for hospice (Gruber et al., 2023), cardiac surgery (Afendulis and Kessler, 2007), skilled nursing therapy (Bowblis and Brunt, 2014; Griffin and Priest, 2024), admissions (Gandhi and Chen, 2023), and discharge decisions (Hackmann et al., 2018). The MUE violations we study can reflect upcoding if only

[3]Evidence of physician hourly misreporting is found in Fang and Gong (2017), although Matsumoto (2020) caution that utilization anomalies alone may not reliably identify fraud.

one procedure is performed but multiple are billed, or unnecessary care if extra procedures are carried out solely to increase reimbursement. A distinctive feature of our setting is that MUE thresholds are medically determined by Medicare and the AMA, yet introduced gradually over time. This staggered rollout allows us to study strategic provider responses and the deterrence effect of these administrative limits.

1. Institutional Background

1.1. Medicare

Medicare is a federal health insurance program that covers more than 60 million Americans, primarily individuals aged 65 and older and younger people with long-term disabilities. Providers typically bill Medicare directly, and reimbursements are determined by standardized fee schedules that specify payments for more than ten thousand procedures. Medicare generally pays about 80% of the approved amount, with beneficiaries responsible for the remaining 20%.

The program's size, complexity, and reimbursement structure create many opportunities for improper billing. Fraud ranges from egregious schemes involving criminal activity to more subtle forms of coding manipulation such as upcoding and padded service intensity. Common schemes include billing for services not rendered, inflating the reported intensity of care, performing medically unnecessary procedures, and accepting kickbacks. Estimates suggest that Medicare fraud costs taxpayers between $100 billion and $300 billion each year (National Health Care Anti-Fraud Association, 2025; U.S. Department of Justice, 2025).

Identifying and prosecuting fraud is challenging. Authorities must establish intent and satisfy the legal requirements of "specificity" and "particularity:" they must point to the exact patient or claim and identify the precise conduct alleged to be fraudulent. Statistical anomalies alone, even when highly improbable, are insufficient under Rule 9(b) of the Federal Rules of Civil Procedure (Li, 2008; Centers for Medicare & Medicaid Services, 2024). These constraints make it difficult to address the more prevalent, subtler forms of coding-

based misconduct and explain the growing interest in administrative tools that aim to block improper payments in real time.

1.2. Medically Unlikely Edits

A central tool in our analysis is the Medically Unlikely Edit program, part of Medicare's National Correct Coding Initiative (NCCI). Introduced for Medicare Part B in 2007, MUEs establish the maximum number of units of a procedure that can plausibly be performed for a patient in a single encounter. These thresholds are determined by CMS in consultation with clinical experts at the American Medical Association and are intended to prevent billing patterns that cannot reflect legitimate medical care.[4] Claim lines that exceed these limits are deemed medically unlikely and are supposed to be automatically denied.

MUEs were initially applied at the claim-line level, meaning the limit was assessed per line on a claim. CMS later expanded the program to include day-of-service MUEs, which cap the total number of units of a procedure that can be billed for a patient on a single day, regardless of how charges are distributed across lines.[5] Not all codes are covered, and CMS updates the list quarterly; most limits are public, while a small subset remains confidential for program-integrity reasons.

MUEs provide an unusually clear structure for identifying implausible billing because they define objective, code-specific upper bounds on medically reasonable service volume. For example, a comprehensive overnight sleep study (HCPCS 95810) can be performed only once per night for a patient. CMS therefore assigns an MUE of one unit for this code, a limit first introduced in 2012. Billing two sleep studies for the same patient on the same day is clinically impossible and constitutes an inherent violation, whether or not the MUE was publicly posted at the time.

In our empirical analysis, we exploit the staggered rollout of these limits. The data

[4] See the CMA description of the MUE program.

[5] See CMS Transmittal 1421.

include both paid and denied claim lines, allowing us to observe the mechanical enforcement of MUE thresholds and to track how providers adjust their billing behavior when new limits become effective. Additional procedural detail on MUE construction, examples, and claim-line versus day-of-service rules appears in Internet Appendix A.

2. Data and Sample

2.1. Medicare Part B Claim-Line Data

Our analysis uses the complete set of Medicare fee-for-service (FFS) Part B Carrier claims from 2012 through 2023. The Carrier file contains professional Part B FFS claims submitted by physicians and other professional suppliers, and it also includes claims from certain free-standing facilities such as independent laboratories, ambulance providers, hospitals, and ambulatory surgical centers. These data originate from the CMS-1500 claim form.[6] We use the Research Identifiable File (RIF) version of the Carrier file, which includes both paid and denied claims and is released annually to approved researchers.

Each Carrier claim contains a header and one or more line items. At the line level, we observe the beneficiary identifier, the date or dates of service, the procedure code (HCPCS or CPT), any modifiers that refine the definition of the service, the number of units billed, and the place of service. The file also reports the submitted charge (the amount billed by the provider), the Medicare-allowed amount, and the actual Medicare payment for each line. Provider identifiers include both the billing and rendering National Provider Identifier (NPI). Using these fields, we construct measures such as total units of service, distinct services per beneficiary-day (counting multiple claim lines with the same code on the same date only once), submitted charges, allowed amounts, Medicare payments, and implied fee-schedule prices by procedure code and place of service.

[6] Additional documentation is available at the ResDAC Carrier file website.

2.2. Final sample

Our initial sample contains the universe of 15 billion claim-line items from 2012 to 2023. We aggregate measures of units, charges, payments, and violations to the provider-procedure-year-quarter level. Thus, in our main specification, the unit of observation is a procedure-level average billed by a specific provider in a specific quarter.

We focus our analysis on the largest providers, hospitals, for two reasons. First, we can links to additional data on financials, leadership, and system ownership, which enables a deeper examination of the determinants of fraud and responses to enforcement. Our final sample includes 6,120 hospitals and their charging behavior for 10,320 procedural codes over the 12 year period from 2012 to 2023.

[Insert Table 1 here]

Table 1 presents summary statistics for our final sample of 157 million hospital-procedure-year-quarter observations. We separately report statistics for claim-line and DoS measures. On average, the amount charged in violation of MUEs is $90 at the claim-line level and $540 at the DoS level. Less than 1% of all claims and charges (regardless of whether a code has an MUE) exceed claim-line limits, while about 2% violate DoS limits. The amount actually paid is substantially smaller than the amount charged, consistent with the general fact that hospital chargemaster rates exceed CMS-negotiated Medicare rates. Notably, incorrect payment rates are roughly five times higher for DoS violations than for claim-line violations. Finally, consistent with DoS limits being introduced later than claim-line MUEs, 80% of observations occur after claim-line MUE introduction compared with 69% for DoS limits.

3. Tracking Fraudulent Medicare Billing Using MUEs

3.1. Concentration of MUE Violations

We begin by examining how MUE violations are distributed across hospitals. The shape of this distribution provides an initial indication of whether the behavior we observe reflects

isolated irregularities or more systematic patterns. Figure 1 plots this distribution. Panel A ranks hospitals by percentiles based on their share of total violations, with the x-axis showing percentiles and the y-axis indicating the fraction of total improper billings accounted for by each group. The distribution is highly skewed: most hospitals contribute negligibly to total violations, while a small subset accounts for the vast majority. Nearly 75% of hospitals exhibit virtually no MUE violations, whereas the top 5% are responsible for roughly 70% of all improper billings. This concentration underscores the presence of a small group of "specialist" hospitals that greatly engage in noncompliant billing behavior and account for a disproportionate share of detected fraud.

[Insert Figure 1 here]

Panel B zooms in on the top 1% of hospitals to further examine this pattern. Even at this level of granularity, concentration remains extreme: fraud is heavily clustered in the top 0.5% of hospitals, which contribute more than twice as much as the average hospital within the top 1%.

Only about 25% of hospitals exhibit any measurable incidence of MUE violations, which raises the question of whether these providers are geographically clustered. Table 2 documents substantial cross-state heterogeneity. Column (1) reports the percentage of submitted claims that violate MUE limits, whereas column (2) reports the same ratio using dollar amounts instead of claim counts. Both measures yield similar magnitudes and an almost identical ranking of states. The five highest-violation states (Vermont, New Hampshire, Alaska, Idaho, and Maine) report violation rates between 4.3% and 5.5%, compared with roughly 1% in lower-violation states (column (1)). Thus, hospitals in these states exhibit rates that are about five times larger than those in the lower end of the distribution. This spatial concentration suggests that enforcement intensity, local regulatory environments, or regional practice norms may play an important role in shaping hospitals' propensity to engage in noncompliant billing behavior.

[Insert Table 2 here]

We next examine the specific procedures in which MUE violations are most common in Table 3. Column (1) list the codes with their frequency in column (2). Three codes account for a large share of improper activity: CPT 88342 (immunohistochemistry staining), HCPCS A9270 (non-covered items or services), and HCPCS G0378 (observation services billed per hour). Violations of 88342 typically involve billing multiple units of the same stain beyond medically plausible limits for a single specimen. For A9270, providers frequently submit repeated units for items that are categorically non-covered. Violations of G0378 arise when hospitals report implausibly high numbers of observation hours for a single patient-day. Additional codes with elevated violation rates include routine, high-volume procedures such as blood tests, colonoscopies, and CT scans.

[Insert Table 3 here]

For the codes listed in Table 3, we conducted a targeted search for anecdotal evidence linking them to allegations of fraud. We classify a code as "high risk" when we find substantial allegations in policy reports or individual lawsuits (column (3)). Most of the codes we flag appear frequently in fraud-related discussions, although a subset shows relatively little external documentation. This contrast highlights the strength of our systematic, data-driven approach: it uncovers patterns of questionable billing even in settings where anecdotal evidence is limited or absent.

3.2. Are MUE Rules Effective at Preventing Improper Billing?

In this section, we examine how billing behavior changes once MUE limits are introduced. If the previously documented violations were primarily driven by legitimate patient needs or other non-behavioral factors, we would expect no significant difference in violation frequency before and after the introduction of MUE limits. We exploit the staggered implementation of MUE limits over time. Specifically, we analyze 2,700 limits introduced

between 2012 and 2024. Figure IA.1 in the Internet Appendix plots the number of new MUE limits introduced by year. The pace of introductions rose rapidly from 2012 to 2018, reflecting the dynamic nature of the healthcare industry, yet new limits continued to appear throughout the sample period. Importantly, after a limit is established, providers may still submit claims that exceed it, but they must include a valid justification explaining why the additional procedures were medically necessary.

Figure 2 displays the unconditional average violation rate, defined as the share of procedure claims exceeding the MUE limit in each year, where year 0 denotes the year in which the MUE limit was introduced. Violations fall sharply, by roughly 95% on average, immediately following implementation, and the decline persists for at least eight years.

[Insert Figure 2 here]

We next turn to regression analysis to account for hospital and procedure characteristics more precisely. The dependent variable is a dummy equal to one if a claim violates an MUE limit. The key explanatory variables are time dummies for periods around the introduction of MUE limits, and we control for hospital size (measured by the number of beds), profitability, and location. We include year-quarter fixed effects to control for unobserved time-specific shocks, since MUE introductions occurred in clusters. Contemporaneous factors such as changes in the economic environment, enforcement intensity, or Medicare reimbursement rules could otherwise induce spurious correlations between the timing of MUE introductions and the observed level of billing violations. Finally, we include HCPCS code fixed effects to absorb average differences across procedures, allowing comparisons of changes in billing behavior for the same procedure over time.

Figure 3 presents the estimated coefficients from the event-time regression, where period 0 corresponds to the introduction of an MUE limit. The shaded area denotes 95% confidence intervals, with standard errors clustered at the hospital level. The figure shows a sharp and immediate decline in violation rates at the time of introduction, followed by a

sustained reduction over the subsequent five quarters. These results confirm that the implementation of MUE limits leads to a significant and lasting decrease in the frequency of improper billing.[7]

[Insert Figure 3 here]

3.3. Are Declines in Violations Driven by Honest Errors?

In this section, we consider a more benign explanation for the observed decline in violations. Specifically, the reduction in improper billing could reflect honest accounting mistakes rather than intentional evasion. To assess this possibility, we conduct three complementary tests that distinguish systematic behavior from inadvertent error.

The first test assesses whether hospitals substitute toward codes that remain unconstrained once an MUE limit is introduced. The logic is straightforward: if violations were largely unintentional, hospitals should reduce questionable billing broadly once an MUE threshold becomes salient for a particular code. By contrast, if the behavior is intentional, constraining one code should lead hospitals to redirect questionable billing toward other codes that do not yet have an assigned limit.

Figure 4 presents the evidence. The red line plots the share of procedure codes not yet assigned an MUE limit. This share declines from about 30% at the start of the sample to roughly 1% by the end, meaning that coverage expands from about 70% to 99% and leaves only a very small (and shrinking) set of uncovered codes where violations remain mechanically possible. The blue line shows the share of claims that violate the remaining (i.e., not yet introduced) MUE limits. As the unconstrained set contracts, the violation rate within it rises sharply: from about 1–2% early in the sample to more than 10% in 2017, before stabilizing near 5%.[8] By the end of the sample, the uncovered codes largely consist

[7] Figure IA.2 in the Internet Appendix shows a comparable pattern when violations are measured in dollar terms and Table IA.1 confirms the result in a standard difference-in-differences setting that considers pre- and post-MUE periods.

[8] Violations do not fall to zero even after limits are introduced because new MUEs continue to be created

of newly introduced procedures not yet assigned a limit.

[Insert Figure 4 here]

This pattern is difficult to reconcile with innocent error. Instead, it points to purposeful substitution. When one avenue for violating limits is closed, hospitals shift activity toward codes that remain outside the enforcement perimeter.

The second test examines whether hospitals that violate MUE limits do so persistently or whether these violations are temporary or sporadic. Figure 5 plots, on the x-axis, the fraction of each hospital's claims that violated MUE limits in a given quarter and, on the y-axis, the same measure for the previous quarter. The scatter plot is color-coded by the frequency density of hospital-quarter pairs: darker red areas represent higher concentrations of observations, while lighter shades indicate lower frequencies. The distribution of points is strongly concentrated along the 45-degree line, indicating that hospitals' violation rates are highly persistent over time.

[Insert Figure 5 here]

We corroborate this visual evidence by estimating persistence directly in Internet Appendix Table IA.2. An AR(1) regression of hospital-level violation rates yields a large and statistically significant autocorrelation of 97.5% (column (1)). Additional regressions with longer lags show that autocorrelation remains positive and statistically significant for up to 20 quarters, declining gradually with the lag length (columns (2) to (6)). This sustained persistence indicates that MUE violations are not random or episodic but instead reflect stable, repeat behavior by a subset of hospitals.

The third test examines the role of financial incentives. If CMS routinely denied claims that exceeded MUE thresholds, violations could plausibly reflect clerical mistakes with no economic payoff. In contrast, if these claims are regularly paid, hospitals face clear incentives to continue submitting them.

over time.

Figure 6 shows that hospitals economically benefit from MUE violations. The blue line plots the share of total Medicare payments attributable to violating claim lines. Before MUE limits become enforceable, this share is substantial, around 0.4 percent of all payments, which aligns closely with the previously documented violation rate. The payment share is slightly lower than the claim share, consistent with Medicare reimbursing only a fraction of submitted charges. The share declines only after the relevant MUE becomes enforceable, matching earlier evidence that MUE limits are generally applied. This pattern is difficult to reconcile with accidental errors and instead suggests that hospitals respond to financial incentives by continuing to submit violating claims as long as those claims generate meaningful payments.

[Insert Figure 6 here]

Taken together, the previous three tests indicate that the patterns we observe are driven by intentional strategies to sustain elevated billing levels despite regulatory limits, rather than by inadvertent errors.

Finally, we conclude this section with one additional cross-sectional comparison. We examine whether hospitals are more likely to exceed MUE limits for procedures with higher reimbursement rates. To do so, we compare violation frequencies across procedures with different payment levels and estimate regressions that relate violation behavior to average reimbursements, including specifications that compare the same procedure code over time.

Table 5 reports regressions of MUE violation rates on average reimbursement amounts. Specifications vary in their inclusion of provider, procedure-code, and quarter fixed effects. Across all models, higher reimbursement amounts are consistently associated with higher violation rates.

[Insert Table 5 here]

3.4. Day-Of-Service MUE violations

We next examine whether hospitals adapt their billing practices when opportunities for excessive billing are curtailed. In particular, we study the introduction of day-of-service (DoS) MUE limits. Unlike the original MUEs, which apply to individual claim lines, DoS limits cap the total number of units of a procedure that can be billed for a patient on a given day, even if the charges are split across multiple lines. These limits were introduced because providers could circumvent claim-line MUEs by dividing a procedure across several lines within the same claim.

Figure 7 plots the share of DoS violations over time. The figure shows that DoS violations remain steady at roughly 1.8% of submitted charges and decline only slightly to about 1.4% after the corresponding limits are introduced. This muted response is consistent with weak or absent enforcement of DoS limits, in sharp contrast to the strong reductions observed for claim-line MUE limits documented above. This pattern also raises the concern that hospitals may shift toward this form of improper billing once claim-line limits are introduced and effectively enforced.

[Insert Figure 7 here]

To explore this mechanism, Figure 8 presents the average rate of claim-line MUE violations (in red) and DoS violations (in blue) in event time, where period 0 corresponds to the introduction of a claim-line MUE limit. Consistent with our earlier findings, claim-line violations fall sharply after the introduction of a limit. At the same time, DoS violations begin to rise almost immediately after implementation. This pattern indicates that hospitals shift from one form of noncompliant billing to another as soon as the original avenue becomes constrained.

[Insert Figure 8 here]

This response is not merely cosmetic. The increase in DoS violations offsets approximately 80% of the reduction in pre-MUE violations achieved by the introduction of claim-

line limits. In other words, hospitals preserve most of the revenue previously generated through claim-line violations by reallocating activity to billing structures that remain permissible under the new rules. This substitution is consistent with hospitals spreading charges for the same procedure across multiple days rather than multiple claim lines within a single day, a strategy that remains technically compliant with DoS rules but defeats the regulatory intent of limiting implausibly high volumes.

Taken together, these patterns provide direct evidence of intentional adaptation. Hospitals do not simply reduce excessive billing when explicit limits are introduced. Instead, they reoptimize to exploit remaining flexibility in the system, shifting to alternative billing methods that are less easily detected. This behavior is once again inconsistent with explanations based on errors or patient complexity and instead reflects purposeful strategies to sustain elevated billing despite stricter regulatory constraints.

4. Conclusion

We exploit a unique time seires of medically grounded, objective upper bounds for the freuqency with which medical actions can be administered to determine the distribution of fraud in medicare, its drivers, the effec tof enforcement and the general equilibrium effects of evasion maneouvers. Based on our analysis of MUE violations in Medicare claims data, we find that the frequency and concentration of violations has decreased over time, likely due to increased monitoring and enforcement efforts. However, we also find evidence that violations are motivated by profit incentives, as violations are concentrated in more lucrative codes. The fraud we document amounts to more than $12 billion over the course of our sample. This estimate represents a lower bound because it captures only clear violations of explicit limits by a subset of actors in a small sub-part of medicare, while many other forms of billing fraud remain.

Overall, our study highlights the importance of taking a comprehensive approach to understanding and addressing fraud in Medicare. By examining patterns at both the indi-

vidual provider level and the organizational level, we are able to provide a more nuanced understanding of the factors that drive MUE violations. Going forward, policymakers and regulators should consider the role of profit incentives and ownership structures in motivating fraud, and explore strategies to reduce these incentives and promote transparency and accountability in the healthcare industry. Ultimately, our findings suggest that efforts to reduce fraud will require a multi-faceted approach, including increased monitoring and enforcement efforts as well as reforms to address underlying systemic issues.

References

Afendulis, Christopher C., and Daniel P. Kessler, 2007, Tradeoffs from integrating diagnosis and treatment in markets for health care, *American Economic Review* 97, 1013–1020.

American Medical Association, 2024, Trends in health care spending Accessed November 2025.

Andreyeva, Elena, Jodi Brown, Sophia Papanastasiou, Katrine V. Seim, and James Traina, 2023, The corporatization of independent hospitals, *Working Paper* .

Arrow, Kenneth J., 1963, Uncertainty and the welfare economics of medical care, *American Economic Review* 53, 941–973.

Bastani, Hamsa, Joel Goh, and Mohsen Bayati, 2019, Evidence of upcoding in pay-for-performance programs, *Management Science* 65, 1042–1060.

Bowblis, John R., and Christopher S. Brunt, 2014, Medicare skilled nursing facility reimbursement and upcoding, *Health Economics* 23, 821–835.

Centers for Medicare & Medicaid Services, 2024, Measuring the value of healthcare anti-fraud efforts Accessed Nov. 2025.

Dafny, Leemore S., 2005, How do hospitals respond to price changes?, *American Economic Review* 95, 1525–1547.

Eliason, Paul J., Mert HeÌli, Patrick Hynes, Anup Malani, and James W. Roberts, 2020, How acquisitions affect firm behavior and performance: Evidence from the dialysis industry, *Working Paper* .

Eliason, Paul J., Riley J. League, Jetson Leder-Luis, Ryan C. McDevitt, and James W. Roberts, 2021, Ambulance taxis: The impact of regulation and litigation on health care fraud, *Working Paper* .

Fang, Hanming, and Qing Gong, 2017, Detecting potential overbilling in medicare reimbursement via hours worked, *American Economic Review* 107, 562–591.

Flasher, Cary, and Maurry Lamboy-Ruiz, 2019, Impact of enforcement on healthcare billing fraud: Evidence from the usa, *Journal of Business Ethics* 157, 177–200.

Gandhi, Ashvin, and Emily Chen, 2023, Picking your patients: Selective admissions in the nursing home industry, *Working Paper* .

Geruso, Michael, and Timothy Layton, 2020, Upcoding: Evidence from medicare on squishy risk adjustment, *Journal of Political Economy* 128, 000–000.

Griffin, John M, Samuel Kruger, and Prateek Mahajan, 2025, Is fraud contagious? social connections and the looting of covid relief programs, *Review of Financial Studies, Forthcoming* .

Griffin, John M., and Alex Priest, 2024, Overbilling and killing? an examination of the skilled nursing industry, *Working Paper* .

Gruber, Jonathan, Mark McClellan, Douglas Staiger, and Xiaotong Xuan, 2023, Dying or lying? for-profit hospices and end of life care, *Working Paper* .

Hackmann, Martin B., Virginia Pohl, and Nicolas R. Ziebarth, 2018, Patient versus provider incentives in long term care, *Working Paper* .

Heese, Jonas, Ranjani Krishnan, and Frank Moers, 2015, Regulator leniency and mispricing in beneficent nonprofits, in *Academy of Management Proceedings*, volume 2015, 1–42.

Howard, David H., and Ian McCarthy, 2020, Deterrence effects of antifraud and abuse enforcement in health care, *Working Paper* .

Howell, Sabrina T., Atul Gupta, Constantine Yannelis, and Abhinav Gupta, 2021, Owner incentives and performance in healthcare: Private equity investment in nursing homes, *Working Paper* .

Jensen, Michael C., and William H. Meckling, 1976, Theory of the firm: Managerial behavior, agency costs and ownership structure, *Journal of Financial Economics* 3, 305–360.

Joiner, Keith A., Jianjing Lin, and Juan Pantano, 2024, Upcoding in medicare: Where does it matter most?, *Health Economics Review* 14, 1–7.

Leder-Luis, Jetson, 2025, Can whistleblowers root out public expenditure fraud? evidence from medicare, *Review of Economics and Statistics* 107, 1015–1033.

Leder-Luis, Jetson, and Anup Malani, 2025, The economics of healthcare fraud, *Working Paper* .

Li, J., 2008, A survey on statistical methods for health care fraud detection, *Health Care Management Science* 11, 275–291, Accessed Nov. 2025.

Liu, Tong, 2022, Bargaining with private equity: Implications for hospital prices and patient welfare, *Working Paper* .

Matsumoto, Brett, 2020, Comment: Detecting potential overbilling in medicare reimbursement via hours worked, *American Economic Review* 110, 3991–4003.

National Health Care Anti-Fraud Association, 2025, The challenge of health care fraud Accessed November 2025.

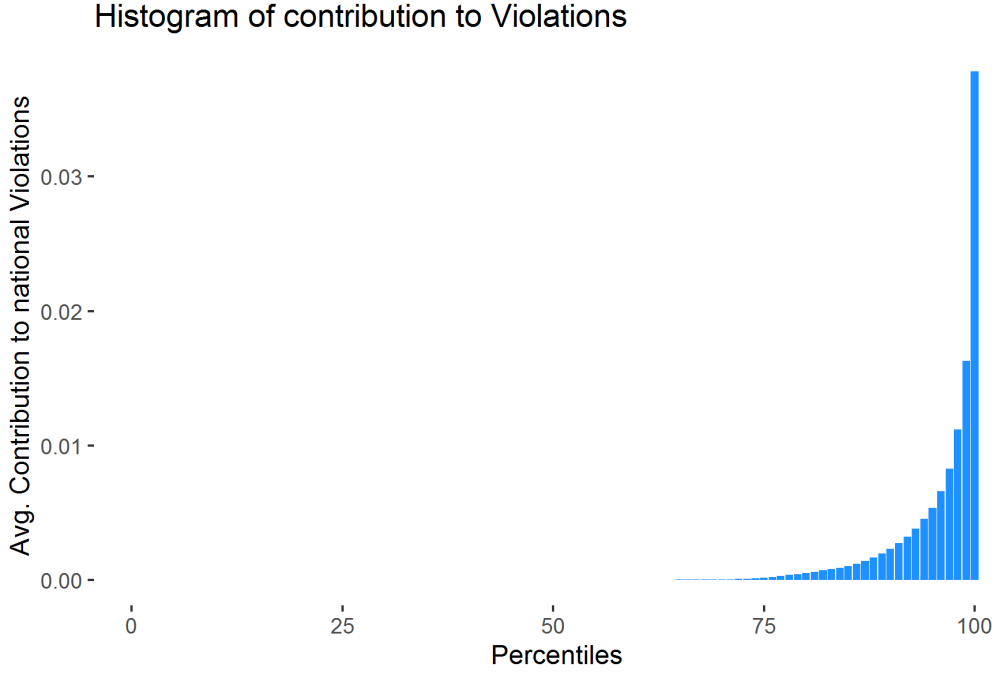
Shekhar, Shubhranshu, Jetson Leder-Luis, and Leman Akoglu, 2023, Unsupervised machine learning for explainable health care fraud detection, *Working Paper* .

Shi, Maggie, 2023, Monitoring for waste: Evidence from medicare audits, *Working Paper* .

Silverman, Elaine, and Jonathan Skinner, 2004, Medicare upcoding and hospital ownership, *Journal of Health Economics* 23, 369–389.

U.S. Department of Justice, 2025, Health care fraudâgenerally Accessed November 2025.

Panel A: Share of Total MUE Violations by Hospital Percentile



Panel B: Share of Total MUE Violations Among Top 1% of Hospitals

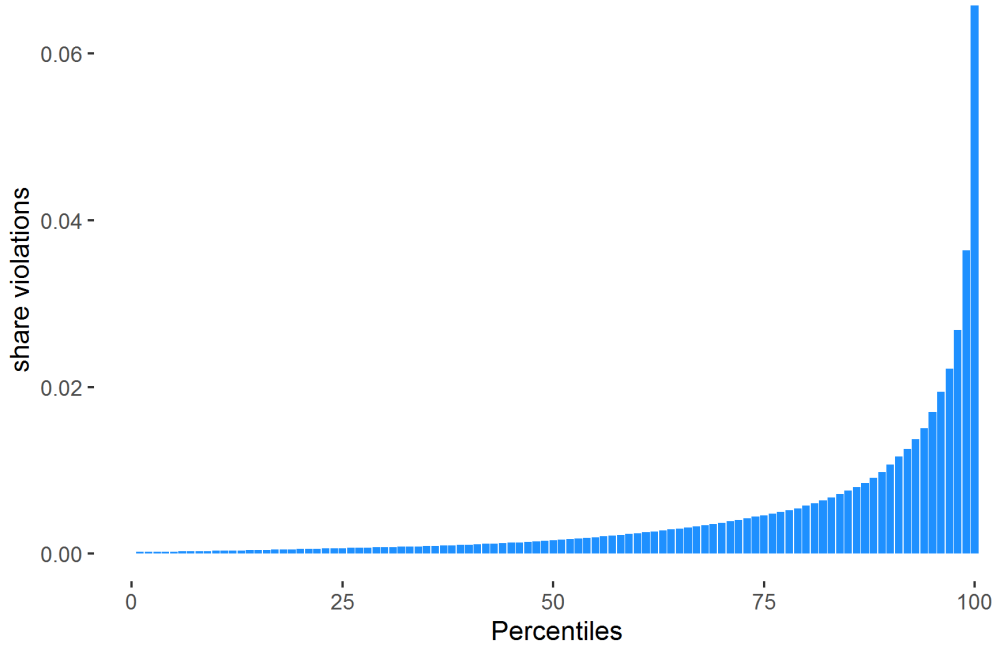


Figure 1.
Distribution of MUE Violations Across Hospitals. This figure displays the distribution of pre-MUE billing violations across hospitals. Panel A ranks hospitals by percentiles according to their share of total violations, with the x-axis representing percentile groups and the y-axis showing the corresponding fraction of total violations. Panel B zooms in on the top 1% of hospitals, presenting the same measure at a finer level of detail to illustrate patterns within the highest-violating providers.

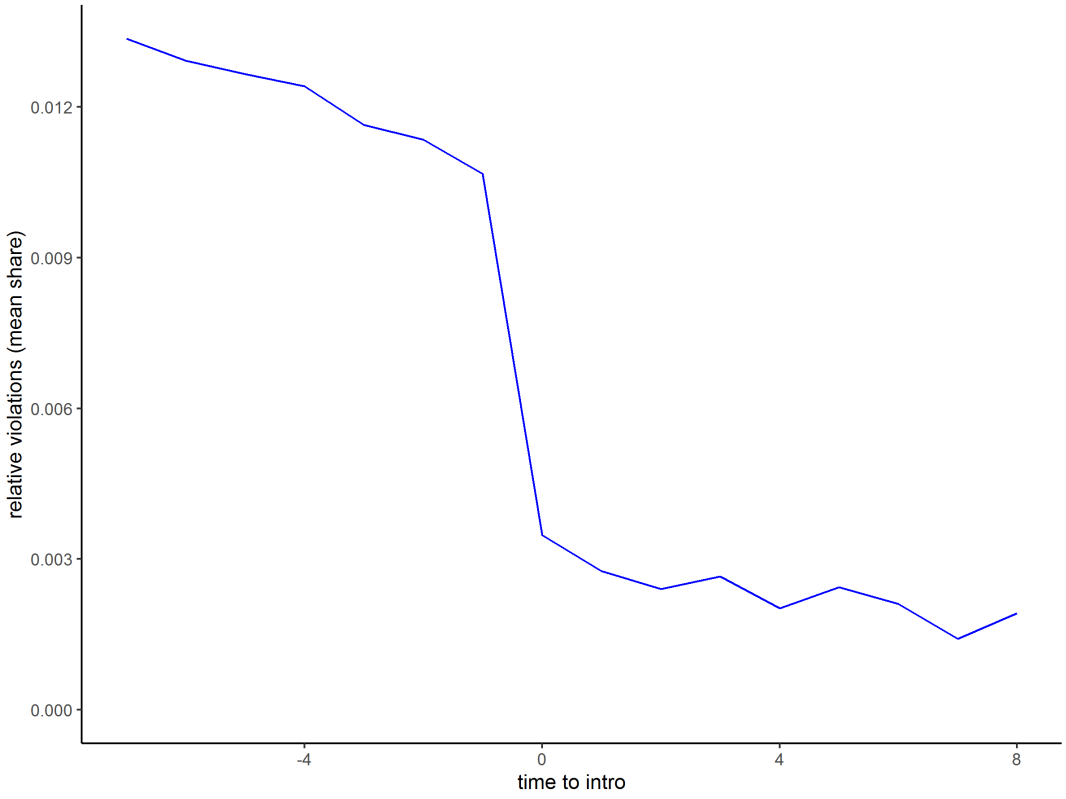


Figure 2.
Average Violation Rates Around MUE Introduction. This figure plots the unconditional average violation rate, defined as the share of procedure claims that exceed the MUE limit, in event time. Year 0 corresponds to the year in which the MUE limit was introduced. Values before and after year 0 show how violation rates evolve around the implementation of the limit.

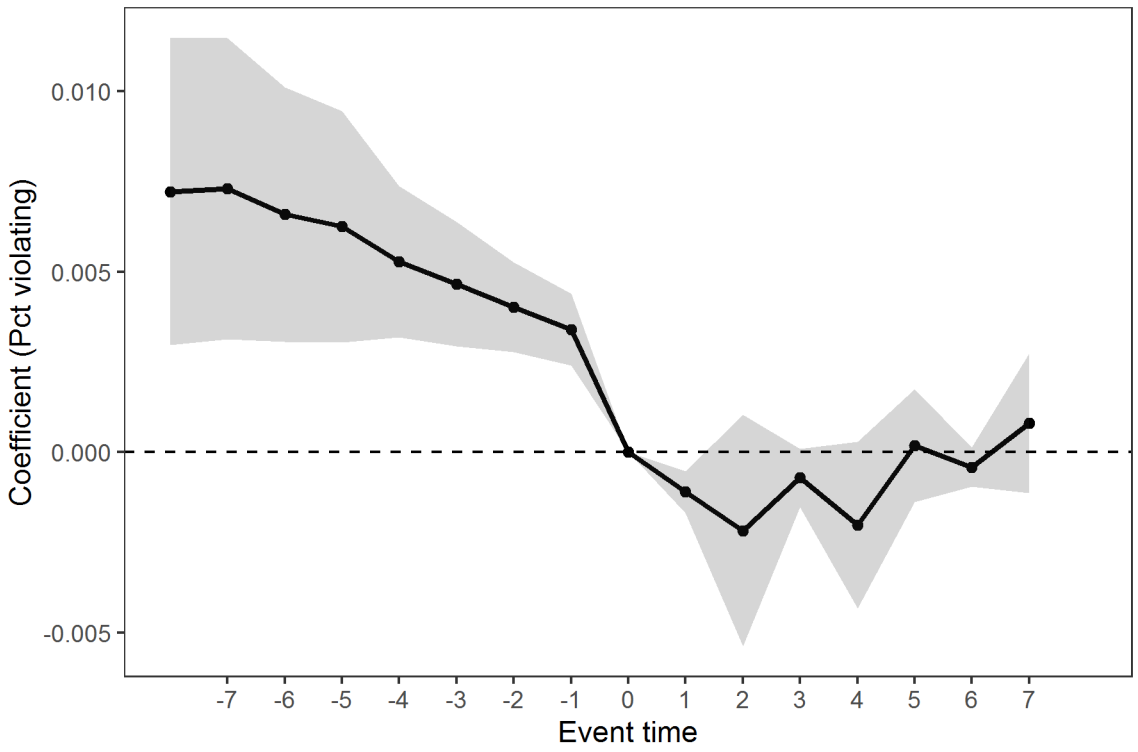


Figure 3.
Event-Time Estimates Around MUE Limit Introductions. This figure plots the estimated coefficients from the event-time regression examining claim-level MUE violations around the introduction of MUE limits. The dependent variable is an indicator equal to one when a claim exceeds the applicable MUE threshold. The key explanatory variables are time dummies for periods relative to the MUE introduction, with period 0 marking the effective date of the limit. Controls include hospital size, profitability, and location. All specifications include year-quarter fixed effects and HCPCS code fixed effects. Shaded areas show 95% confidence intervals, with standard errors clustered at the hospital level.

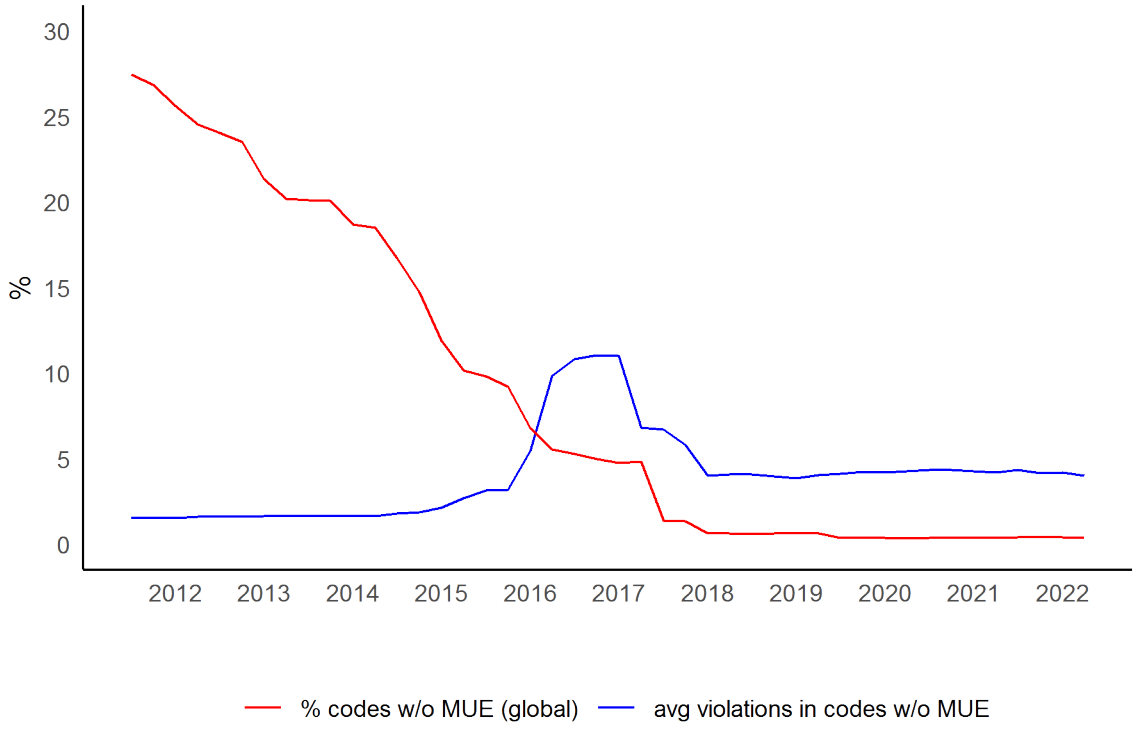


Figure 4.
Coverage of MUE Limits and Violation Rates Among Uncovered Codes. ed line shows the share of HCPCS procedure codes that do not yet have an assigned Medically Unlikely Edit (MUE) limit. The blue line shows the share of claims that exceed the MUE thresholds for the subset of procedure codes that remain uncovered at each point in time.

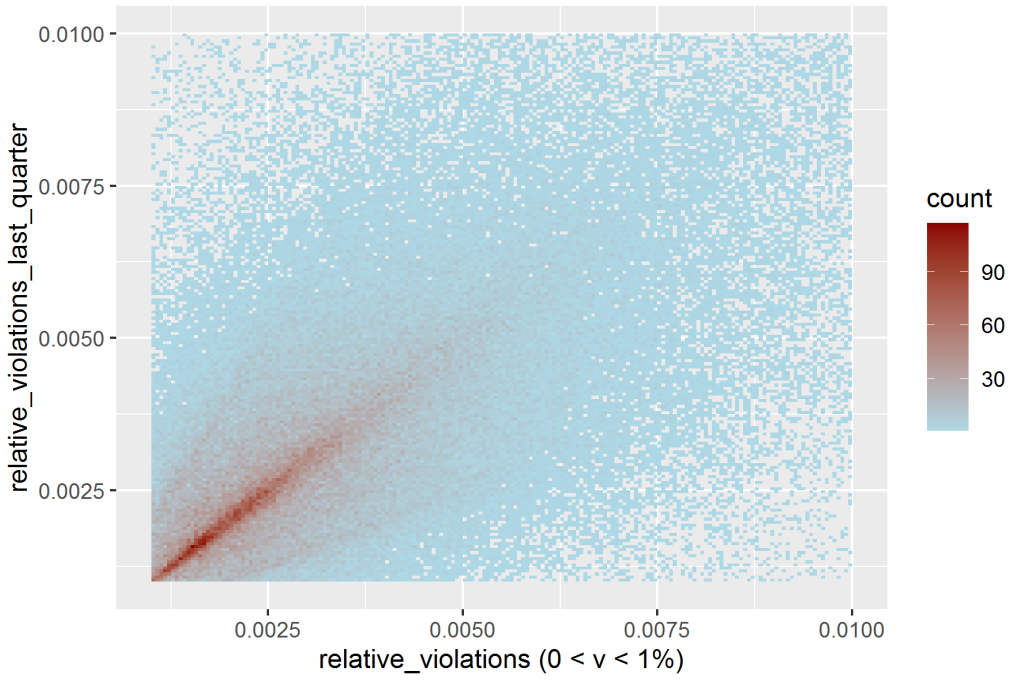


Figure 5.
Quarter-to-Quarter Persistence of Hospital MUE Violation Rates. This figure plots, for each hospital and quarter, the fraction of submitted claims that violated MUE limits (x-axis) against the same measure in the preceding quarter (y-axis). Each point represents a hospital-quarter pair, and shading indicates the frequency density of observations, with darker colors denoting higher concentrations. The figure illustrates how violation rates relate across adjacent quarters.

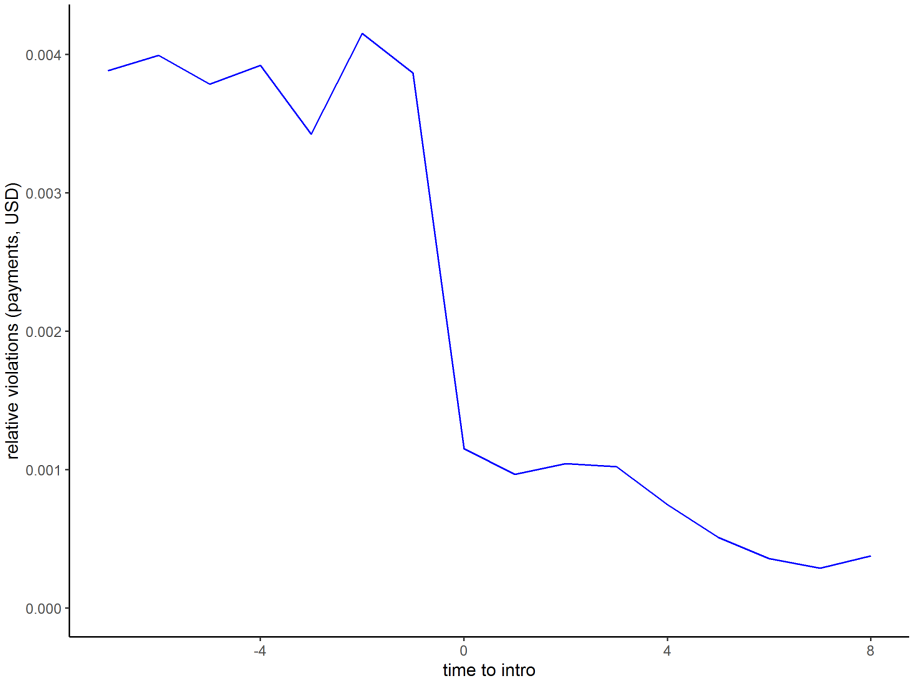


Figure 6.
Payments Associated with MUE-Violating Claims. This figure plots the share of total Medicare payments attributable to claim lines that exceed Medically Unlikely Edit (MUE) limits. The blue line reports the fraction of all Medicare payments linked to violating claims in each period.

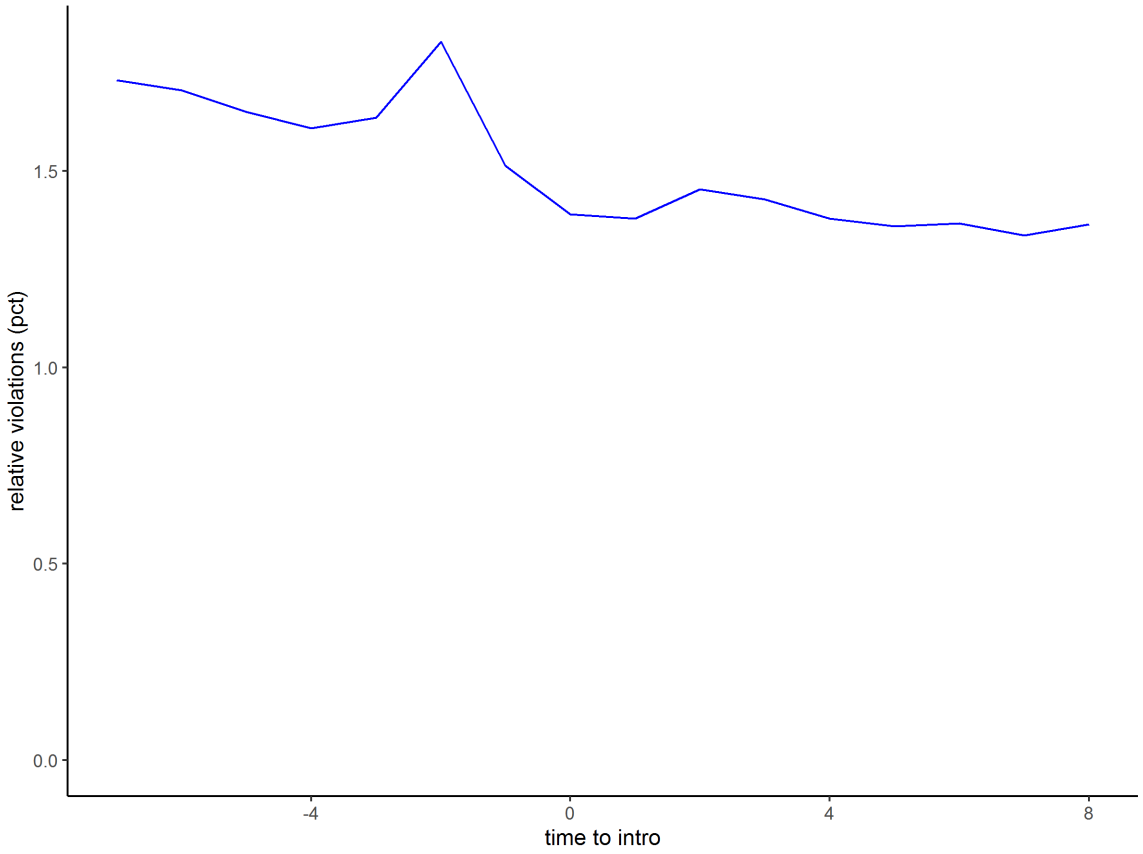


Figure 7. Day-of-Service (DoS) Violation Rates in Event Time Around DoS MUE Introduction. This figure plots the share of submitted charges that violate day-of-service (DoS) MUE limits in event time, where period 0 corresponds to the quarter in which a DoS limit is introduced.

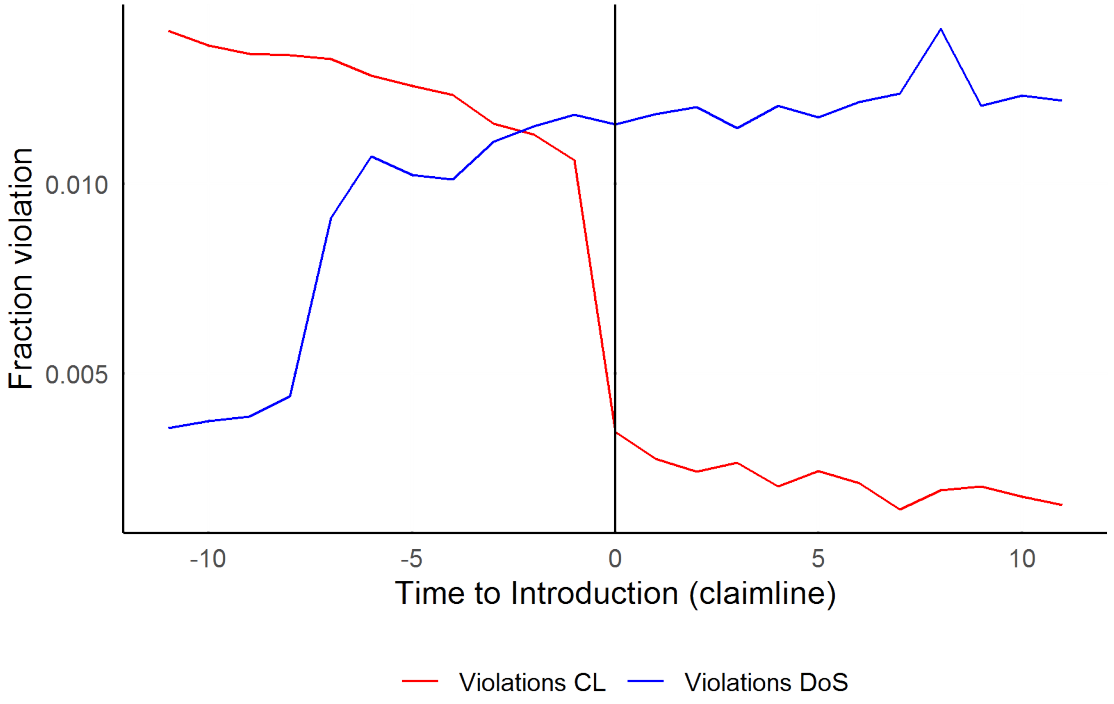


Figure 8. Claim-Line and Day-of-Service Violation Rates Around MUE Introduction. This figure plots the average share of claim-line MUE violations (red) and day-of-service (DoS) violations (blue) in event time, where period 0 denotes the quarter in which a claim-line MUE limit is introduced.

Table 1
Summary statistics
This table reports summary statistics on variables at the hospital-procedure-year-quarter level.

Variable	N	Mean	SD	P25	P50	P75
Claimline						
$ charged in violations	157,397,164	87.57	4492.52	0	0	0
Violations as fraction of claims	157,397,164	0	0.05	0	0	0
$ charged in violations, relative to all charges	156,710,268	0	0.05	0	0	0
$ paid by CMS in violations, relative to all charges	157,397,164	2.70	714.18	0	0	0
1(Post MUE introduction)	157,397,164	0.80	0.40	1.00	1.00	1.00
Day of Service						
$ charged in violations	157,415,954	541.08	26934.17	0	0	0
Violations as fraction of claims	157,415,954	0.02	0.12	0	0	0
$ charged in violations, relative to all charges	156,729,103	0.02	0.12	0	0	0
$ paid by CMS in violations, relative to all charges	157,415,954	62.71	4989.76	0	0	0
1(Post MUE introduction)	157,415,954	0.69	0.46	0	1.00	1.00

Table 2
State-Level Incidence of MUE Violations
This table reports the incidence of Medically Unlikely Edit (MUE) violations across U.S. states. Column (1) shows the percentage of submitted Medicare Part B claim lines that exceed the applicable MUE limit. Column (2) reports the corresponding share of total Medicare claimed amounts associated with violating claim lines.

State	MUE violations (%)	MUE violations (USD %)
Vermont	5.49	5.75
New Hampshire	4.92	5.20
Alaska	4.52	4.81
Idaho	4.36	4.66
Maine	4.29	4.58
Montana	4.22	4.50
Oregon	4.10	4.39
Iowa	3.83	4.10
Wyoming	3.74	4.04
Wisconsin	3.74	3.99
North Dakota	3.51	3.75
Minnesota	3.43	3.67
Washington	2.88	3.13
Nebraska	2.11	2.34
South Dakota	2.13	2.33
Colorado	1.96	2.24
Arizona	1.85	2.16
District of Columbia	1.61	2.03
Michigan	1.70	1.96
Mississippi	1.68	1.92
Kansas	1.72	1.91
Hawaii	1.68	1.88
West Virginia	1.61	1.83
Illinois	1.58	1.83
Utah	1.58	1.79
Missouri	1.50	1.74
Arkansas	1.46	1.72
New Mexico	1.44	1.68
California	1.37	1.65
Massachusetts	1.33	1.58
New York	1.32	1.57
Texas	1.26	1.53
Nevada	1.27	1.52
Maryland	1.16	1.47
Louisiana	1.20	1.46
Ohio	1.24	1.46
North Carolina	1.21	1.45
Indiana	1.21	1.45
Oklahoma	1.21	1.44
South Carolina	1.18	1.40
Rhode Island	1.15	1.39
Pennsylvania	1.16	1.39
Kentucky	1.15	1.39
New Jersey	1.09	1.39
Georgia	1.09	1.33
Connecticut	1.09	1.30
Virginia	1.04	1.26
Delaware	0.96	1.21
Alabama	0.98	1.21
Puerto Rico	1.00	1.18
Florida	0.94	1.17
Tennessee	0.90	1.09

Table 3
Most Frequently Violated HCPCS/CPT Codes
This table lists the procedure codes associated with the highest number of MUE violations in our data. Column (1) reports the share of total MUE violations attributable to each code and its frequency (column (2)). Column (3) classifies codes as "High Risk" when independent searches reveal anecdotal evidence of fraud allegations in policy reports or individual lawsuits.

Code	Name	Frequency	High risk?
88342	Immunohistochemistry, each antibody	63374	Yes
A9270	Non-covered item or service	61191	Yes
G0378	Hospital observation service, per hour	49729	Yes
88185	Flow cytometry, each additional marker	48232	Yes
82962	Glucose, blood by home monitoring device	24339	No
J9299	Injection, nivolumab (Opdivo) 1 mg	19898	Yes
43239	EGD (esophagogastroduodenoscopy) with biopsy	12223	No
45380	Colonoscopy with biopsy	12093	No
J9271	Injection, pembrolizumab (Keytruda) 1 mg	10323	Yes
45385	Colonoscopy with removal of polyp by snare	9836	Yes
L8699	Prosthetic implant, not otherwise specified	8950	No
77386	IMRT delivery (intensity modulated radiation therapy)	8648	Yes
99291	Critical care, 30-74 minutes	8601	Yes
C1713	Anchor/screw for bone-to-bone or soft tissue-to-bone	8409	Yes
93306	Transthoracic echocardiography (complete)	8129	Yes
80053	Comprehensive metabolic panel	7771	Yes
88341	Immunohistochemistry, each additional antibody	7300	Yes
72125	CT scan of cervical spine without contrast	7296	No
77300	Basic radiation dosimetry calculation	6700	Yes
Q9967	Low-osmolar contrast material, per mL	6667	Yes

Table 4
Elasticity of DoS violations with respect to claimine violations
This table reports regressions of code-quarter average MUE-DoS violation rates (in logs) on average claimline violation rates (in logs). Heteroskedasticity robust standard errors clustered on the procedure level are reported in parenthesis. *,**,*** denote statistical significance on the 10%, 5% and 1% level, respectively.

	log(Violations DoS (fraction))
	(1)
log(Violations CL (fraction))	-0.17^{***}
	(0.02)
Adj. R^2	0.38
Num. obs.	92

Table 5
Reimbursement Levels and MUE Violation Rates
This table reports regressions of MUE violation rates on average reimbursement amounts. Specifications vary in the inclusion of provider, procedure-code, and quarter fixed effects. Heteroskedasticity robust standard errors clustered on the procedure level are reported in parenthesis. *,**,*** denote statistical significance on the 10%, 5% and 1% level, respectively.

	(1)	(2)	(3)
Avg_payment_per_procedure	0.01***	0.01***	0.01***
	(0.00)	(0.00)	(0.00)
Num. obs.	157047820	157047820	157047820
Adj. R^2 (full model)	0.40	0.40	0.40
Provider FE	6180	6180	
Procedure FE	16706	16706	16706
Year-quarter FE		48	
Provider × Year-quarter FE			259354

Internet Appendix for "How Effective is Ex Ante Enforcement? Evidence From 15 Billion Medicare Claims"

A. Additional Details on MUE Rules and Billing Mechanics

This appendix provides supplementary institutional detail on how Medically Unlikely Edits are defined and implemented, expanding on the brief overview in Section 1.2.

A.1. Structure of MUE Thresholds

Medicare outpatient claims are submitted at the claim-line level, where each line corresponds to one HCPCS/CPT code and the number of units billed for that procedure. Within the National Correct Coding Initiative (NCCI), the Medically Unlikely Edits (MUE) program establishes medically plausible upper limits on these units. These thresholds are determined jointly by clinical experts at CMS and the American Medical Association and are intended to capture the maximum feasible quantity of a service that can be legitimately provided to a patient in one encounter. MUEs take two forms:

1. **Claim-line MUEs**, which limit the units of a code within a single claim line.
2. **Day-of-service (DoS) MUEs**, which cap the total number of units of a code that can be billed for a patient on a given date, regardless of how many claim lines are used.

Claim-line MUEs were introduced beginning in 2007, and DoS limits were added later for a subset of codes. CMS updates both sets of limits quarterly.

Providers are generally required to report all units of a procedure on one line. Splitting a service across multiple lines is allowed only in narrow circumstances and requires modifiers that document the clinical justification.[1]

A.2. Illustrative Example: Sleep Study Code 95810

To illustrate how MUEs operate, consider HCPCS code 95810, which describes an overnight polysomnography sleep study attended by a technologist. This procedure occupies an entire night and cannot be performed twice on the same day for the same patient. In 2012, CMS introduced an MUE of one unit for this code, justified by "the nature of the procedure."

Submitting two units on a single claim line should therefore trigger an automatic denial. Even prior to the MUE introduction, billing more than one unit was incompatible with medically plausible practice. The MUE simply makes enforcement explicit and automatic.

[1] See CMS NCCI FAQ library: link.

A.3. Claim Splitting as an Evasion Strategy

After MUE limits are introduced, some hospitals attempt to obtain payment by splitting units across multiple claim lines instead of placing all units on a single line. For example, instead of billing two units of 95810 on one line, they submit two separate lines each with one unit. Although this practice is generally not permitted under CMS rules, it can temporarily evade the automatic enforcement of claim-line MUEs.

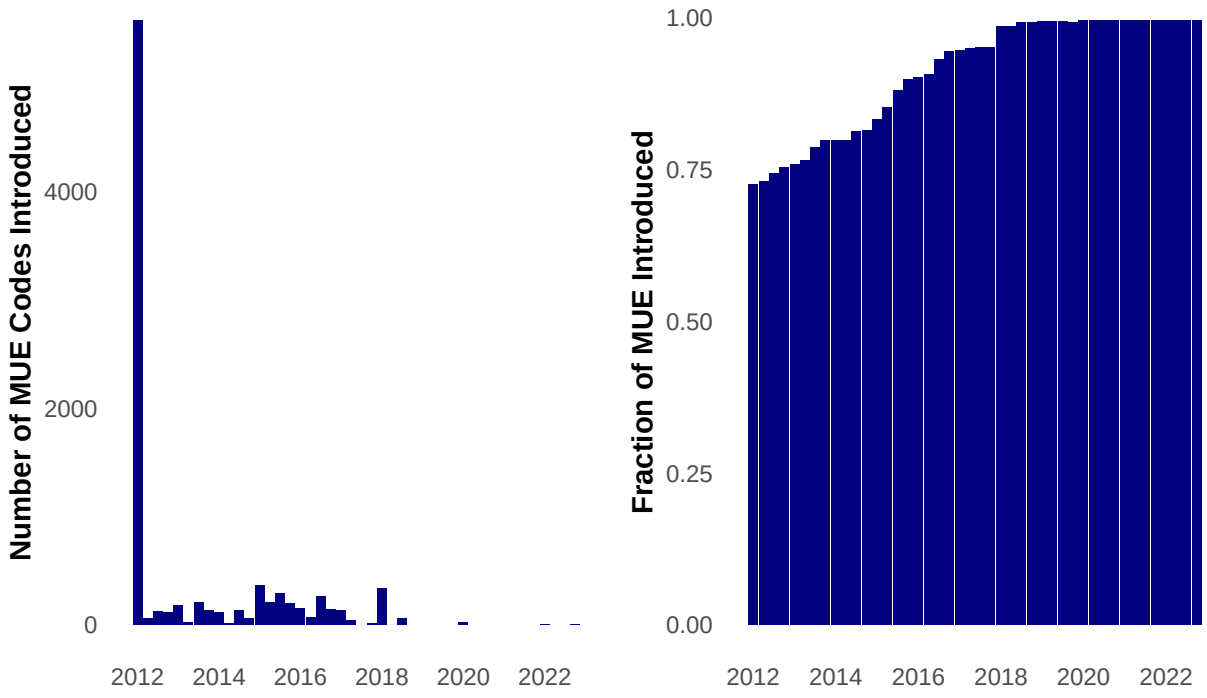
Splitting is only allowed under narrow circumstances, for example, if one physician repeats a study performed earlier that day, and in those cases the claim must include appropriate modifier codes. Importantly, even when splitting occurs, DoS MUEs cap the total number of units in a day and require denial of all units that exceed this cap. A concrete illustration of this structure appears in Internet Appendix Figure IA.3.

A.4. DoS MUE Violations and Enforcement Gaps

Not all codes have both claim-line and DoS MUEs, and DoS thresholds are introduced later, as shown by the summary statistics on time to introduction. We separately track violations of DoS limits and document cases where hospitals exceed the daily cap without violating claim-line limits. These violations play an important role in the strategic substitution documented in the main text.

B. Supplementary Figures and Tables

Figure IA.1.
Annual Introduction of New MUE Limits and saturation. The left panel in this figure plots the number of Medically Unlikely Edit (MUE) limits introduced each year between 2012 and 2024. The right panel plots the cumulative saturation of existing procedure codes with claim line level MUE limitss.



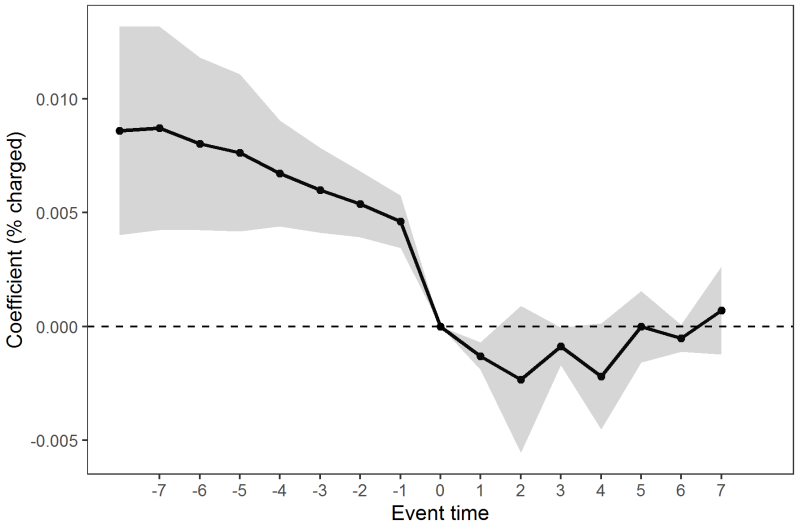


Figure IA.2.
Average Dollar-Based Violation Rate Around MUE Introduction. This figure plots the unconditional average violation rate measured in dollar terms, where the numerator is the total dollar amount of claims that exceed the MUE threshold and the denominator is total dollar spending on the relevant codes in each year. Year 0 corresponds to the year in which the MUE limit was introduced.

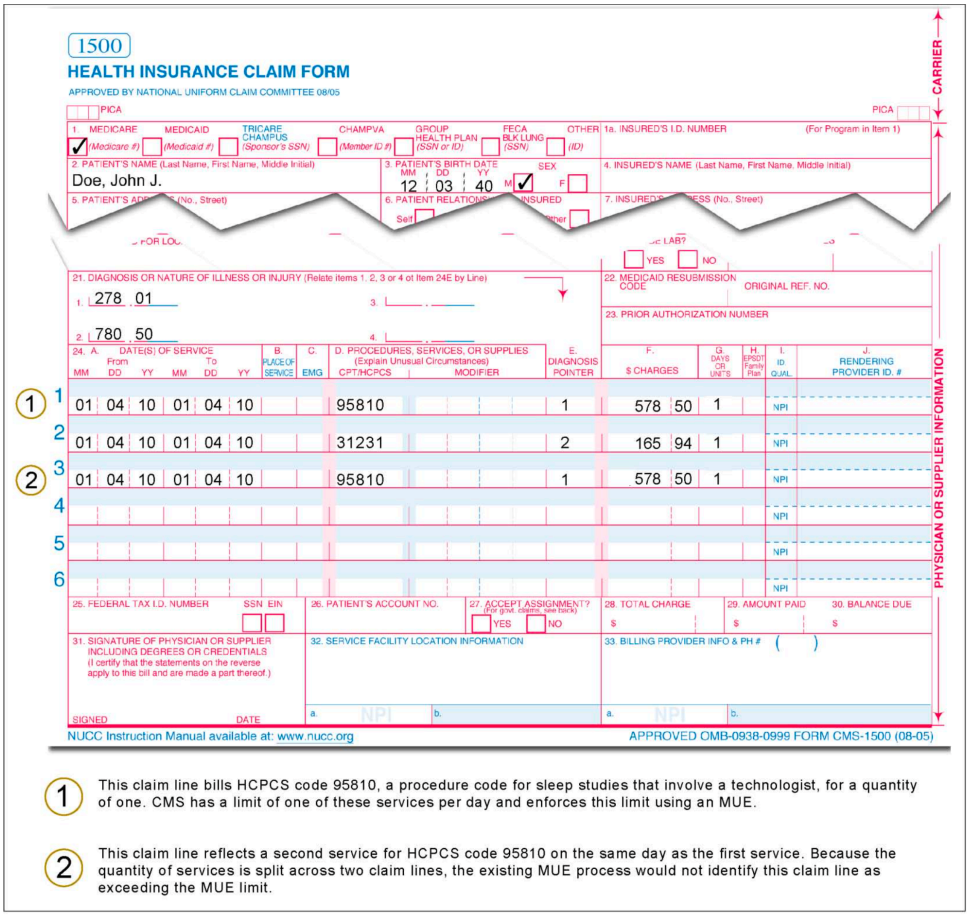


Figure IA.3.

Example claim. Illustration of how units of a single HCPCS code can appear across multiple claim lines. In this example, units that would exceed the claim-line MUE if billed together are instead split across two lines. Although each line individually falls below the claim-line threshold, the aggregate quantity exceeds the permissible day-of-service limit.

Table IA.1

Dynamics relative USD charged around introduction (claimline)

This table presents results from a difference in differences regression of the fraction of violating charges on treatment × post, the interaction of an indicator taking value 1 for codes with an MUE limit introduced at some point and an indicator for years after introduction. Heteroskedasticity robust standard errors clustered on the procedure level are reported in parenthesis. *,**,*** denote statistical significance on the 10%, 5% and 1% level, respectively.

	(1)	(2)	(3)	(4)
Treatment × Post	-0.01^{***}	-0.01^{***}	-0.01^{***}	-0.01^{***}
	(0.00)	(0.00)	(0.00)	(0.00)
Num. obs.	37043142	37043142	37043142	37043142
Adj. R^2 (full model)	0.35	0.36	0.37	0.38
Num. groups: hcpcs_cd	10320	10320	10320	10320
Num. groups: yearq2		48	48	
Num. groups: prvdr_num			6120	
Num. groups: prvdr_yearq				233022

Table IA.2
Persistence of MUE Violation Rates Across Hospitals
This table reports autocorrelation estimates from regressions of hospital-level MUE violation rates on their lagged values. Heteroskedasticity robust standard errors clustered on the procedure level are reported in parenthesis. *,**,*** denote statistical significance on the 10%, 5% and 1% level, respectively.

Dependent Variable:	total_violations					
	(1)	(2)	(3)	(4)	(5)	(6)
L1(total violations)	0.9748***					
	(0.0012)					
L4(total violations)		0.9091***				
		(0.0064)				
L8(total violations)			0.8287***			
			(0.0124)			
L12(total violations)				0.7508***		
				(0.0177)		
L16(total violations)					0.6875***	
					(0.0219)	
L20(total violations)						0.6605***
						(0.0283)
Fixed-effects						
yearq2	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	1,204,661	1,057,334	872,308	698,149	534,208	379,222
R^2	0.95180	0.82769	0.70016	0.59639	0.51186	0.45004