Fragmented Insurance and Billing Frictions: Understanding Denied Health Insurance Claims*

Riley League† Mark Shepard‡ Myles Wagner§

May 2026

Billing and insurance-related activities are a major contributor to the high administrative costs of the US health care system. Yet the economics of claim denials—one of the most visible manifestations of billing frictions—remains poorly understood. We argue that the prevalence of denials presents an economic puzzle and study why they persist in equilibrium. Using data from the Massachusetts All-Payer Claims Database, we document that claim denials are widespread across insurers, occur frequently for services that an insurer predictably denies, and are especially common for routine, low-cost services. Using a simple model of insurer billing, we argue that insurer *fragmentation*, combined with limited provider ability to tailor billing strategies across insurers, can explain these otherwise puzzling patterns. Consistent with this mechanism, denial rates rise when patients switch insurers and vary substantially across insurers for the same services, suggesting providers do not finely tailor billing across payers. We also show that denial rates are lower for providers with more billing experience and for vertically integrated provider–insurer pairs. Together, our results highlight how fragmented insurance systems and heterogeneous coverage rules can generate administrative inefficiencies and contribute to the high costs of the US health care system.

*Riley League thanks the National Institute on Aging, grant number T32-AG000186, for its support.

†Gies College of Business, University of Illinois Urbana-Champaign and NBER

‡Harvard University and NBER

§The Ohio State University

1 Introduction

Administrative costs are a defining feature of the US health care system. A large and growing literature documents that health care spending in the United States substantially exceeds that of other high-income countries, and that administrative costs account for a meaningful share of this gap. Estimates suggest that administrative costs represent roughly 15–30 percent of total US health care spending, on the order of one trillion dollars annually, and explain a substantial fraction of the difference between US spending and that of peer countries (Pozen and Cutler, 2010; Jiwani et al., 2014). Much of this administrative burden arises from billing and insurance-related activities, including claims management, clinical documentation, coding, and utilization review (Jiwani et al., 2014). Despite their scale, these frictions remain poorly understood from an economic perspective, in part because they are diffuse, difficult to measure, and deeply embedded in the institutional structure of private insurance markets.

We focus on denied health insurance claims, which occur when an insurer declines to pay for a service after it has been rendered by a provider. Denied health insurance claims are among the most visible, common, and controversial manifestations of these administrative frictions. Virtually all claims submitted to insurers face the threat of a denial, and nationally between 12–20% of claims are actually denied (Optum, 2024; Lo et al., 2025; Vabson et al., 2025). Claim denials are costly, both because their appeal involves substantial paperwork costs (Dunn et al., 2024) but also because providers often take costly actions, like hiring staff or adopting billing technology, to preemptively avoid them (League, 2024; Becker's Hospital Review, 2024; Forbes Technology Council, 2024). Being common and costly, they have garnered significant public attention, almost all of it negative, with large majorities across political affiliation, income, and insurance type citing denials as a major problem (Kaiser Family Foundation, 2025).

The high prevalence of claim denials represents an economic puzzle. First, in the absence of frictions, providers should internalize billing rules and avoid submitting claims that will result in a denial. Second, since denials are costly, Coasean bargaining logic suggests insurers and providers have strong incentives to design billing processes that minimize unnecessary frictions, which could preserve net transfers from insurers to providers with lower administrative costs (Coase, 1960).

Nonetheless, claim denials are extremely common. In this paper, we ask why this is, economically. We provide new evidence on the nature and origins of denied health insurance claims in the United States. We use detailed administrative data from the Massachusetts All-Payer Claims Database (MA-APCD) to document key facts about claim denials across nearly the entire private insurance landscape. Our central contribution is to show that denials are closely related to *fragmentation* in health insurance markets: the presence of many insurers pursuing distinct and uncoordinated approaches to claims review. We document substantial heterogeneity in how

insurers deny claims for the same services, as well as sharp increases in denials when individuals switch plans or hold overlapping coverage. Taken together, these patterns point to a system characterized by fragmented enrollment and fragmented rules, in which a lack of coordination between insurers and providers generates large administrative frictions.

We begin by documenting several basic descriptive facts about denied claims. Denials are widespread across all major segments of the private insurance market, including employer-sponsored insurance, the individual market, Medicaid managed care, and Medicare Advantage. Averaging across years, roughly 13–15 percent of claim lines are denied, a rate that is remarkably stable over our 2013–2020 data period. These estimates are broadly consistent with national benchmarks reported in industry data (Optum, 2024). Denials are not concentrated in a particular insurer type or time period, suggesting that they reflect a persistent feature of the billing system rather than transitory policy changes or short-run cost shocks.

Next, we examine which services are most likely to be denied. In contrast with a simple cost-containment narrative, denials are substantially more common for low-price and routine services than for expensive procedures. Denial rates decline sharply with the average price of a procedure, even after excluding extremely low-priced services. Many of the most frequently denied services are common features of routine care, such as blood draws, laboratory tests, and basic diagnostic procedures. We also find little evidence that denials are targeted toward sicker or higher-spending patients, even in individual health insurance markets where risk selection incentives are plausibly strongest. These patterns are difficult to reconcile with the view that denials primarily serve to discourage high-cost utilization or for insurers to select healthy patients.

Third, we present evidence that many denials are for services that are nearly certain to be denied *ex ante*. To show this, we compute insurer-by-procedure denial rates and examine the distribution of these rates across all submitted claims. We find that a non-trivial share of claims—approximately 3 percent—are submitted to insurers that deny over 80 percent of claims for that procedure. These predictably denied claims account for a hugely disproportionate share of all denials: 28 percent of denied claim lines involve insurer-procedure combinations with denial rates above 80 percent.

With these stylized facts in hand, we turn to explaining them using an economic model of claim denials. We develop a simple but flexible auditing model where claim denials arise when insurers catch providers attempting to overbill, when providers bill for non-covered care, or when there is a mismatch between provider and insurer billing styles. When there is a single provider and single insurer, the model cannot generate the patterns we see empirically. However, introducing multiple insurers with different coverage rules and billing styles, along with an inability on the part of providers to perfectly tailor their billing to each insurer, can match the stylized facts we observe in the data.

With theoretical justification that insurer fragmentation can generate the patterns we observe

in the data, we then present evidence of this fragmentation generating claim denials. First, we link claim denials to fragmentation in insurance *enrollment*: the fact that consumers regularly switch insurance plans, requiring providers to adapt claims submissions to a different insurer and its billing practices. Using an event-study design, we show that denial rates spike sharply when individuals switch insurers. Among "clean" switchers—individuals who change plans without gaps in coverage or overlapping enrollment—denials increase by approximately 5 percentage points at the time of the switch, a roughly 35% increase relative to baseline rates. Much of this spike reflects a simple error: the provider submits claims to the wrong (old) insurer after the patient switches plans. This reflects a simple type of miscommunication in a fragmented insurance system without a central database for tracking enrollment. Additionally, denial rates remain elevated for many months afterward, suggesting the new insurer may require additional documentation to verify medical necessity. Switching insurers is common, with roughly 10 percent of individuals changing insurers in a given year, suggesting the aggregate importance of this mechanism. We also find higher denial rates among individuals with simultaneous dual coverage in multiple plans (e.g., a plan from both their own and their spouse's employer), consistent with difficulties coordinating benefits across insurers.

Next, we provide evidence that fragmentation in insurer *coverage rules* generates claim denials. Insurers differ in whether particular services are covered at all, whether they are paid separately or bundled into broader claims for a visit, and whether they are subject to network restrictions or service limits. Providers may not fully understand insurer-specific rules and submit claims in a "coarse" manner, resulting in predictable denials.

To study the consequences of this mismatch, we analyze variation in insurer-specific denial rates for a given procedure. For many routine services—such as blood draws, pulse oximetry tests, physical therapy services, and refractive eye exams—the distribution of denial rates across insurers is strongly *bimodal*: some insurers almost always deny the service, while others rarely do. These bimodal patterns are especially pronounced for services with high overall denial rates and account for a substantial share of all denied claim lines. In contrast, services with clear and uniform coverage rules, such as standard physician office visits, exhibit low denial rates and little cross-insurer dispersion.

We implement a simple method to quantify the prevalence of "bimodal" denial procedures, for which denial rates differ widely across insurers. We find that bimodal procedures are common, representing 55% of all claims lines. They comprise a disproportionate share of denials, representing 69% of all denied claims. This suggests the importance of fragmented rules across insurers for understanding broad denial patterns.

As a final piece of evidence for the importance of insurer fragmentation, we show that vertically integrated health systems experience fewer denials when system providers bill the vertically integrated insurer. This is consistent with vertical integration solving billing frictions through

administrative coordination, weakening the incentives to bill and deny aggressively, and increased volume familiarizing providers with billing rules. We also provide independent evidence of this last channel, with the denial rate being decreasing in the number of claims a provider sends to an insurer.

Taken together, our findings suggest that denied claims are not only a tool for selectively targeting high-cost care or high-risk patients. Instead, they appear connected to a fragmented insurance system in which multiple insurers apply heterogeneous billing rules to a common set of provider claims.

Related Literature This paper contributes to three distinct literatures. First, we contribute to the literature on administrative costs in US health care by providing micro-level evidence on a specific, measurable source of administrative waste: denied claims. A large body of work documents that the US spends far more on health care administration than peer countries, attributing much of this gap to the complexity of the multi-payer system (Woolhandler et al., 2003; Himmelstein et al., 2014; Cutler and Ly, 2011; Jiwani et al., 2014). Complementary studies quantify the costs that providers bear in navigating insurer-specific rules for claims submission, prior authorization, and credentialing (Casalino et al., 2009; Tseng et al., 2018; Cutler et al., 2012). A growing economics literature examines claim denials and related billing frictions in public insurance programs, documenting that denials impose large financial costs on providers (Dunn et al., 2024), that providers adapt billing practices in response to denial intensity without reducing insurer spending (League, 2024), and that utilization management tools such as prior authorization meaningfully shape provider behavior (Brot-Goldberg et al., 2023; Shi, 2024; Gandhi and Shi, 2025; Eliason et al., 2025). We complement this work by analyzing denials across nearly the entire private insurance landscape and by showing that denial patterns are inconsistent with simple cost-containment explanations: denials are concentrated among low-price, routine services and are not disproportionately targeted at high-cost enrollees. These findings suggest that administrative frictions arise at least partly from the institutional structure of insurance markets rather than simply from clinical discretion or strategic rationing.

Second, we contribute to the theory of contracting and auditing by extending standard costly state verification models (e.g., Mookherjee and Png, 1989; Crocker and Morgan, 1998) to a setting with multiple insurers. In the canonical single-insurer framework, the insurer audits claims at a cost, and the equilibrium features low denial rates, no denials for inexpensive procedures, and no submission of claims for non-covered services. We show that introducing multiple insurers with heterogeneous coverage rules—combined with providers' limited ability to tailor billing strategies to each payer—fundamentally alters these predictions. We characterize an equilibrium in which providers adopt coarse billing strategies that are optimal on average but generate predictable denials from insurers whose rules differ from the norm. The resulting model produces high equi-

librium denial rates, a negative relationship between procedure price and denial probability, and bimodal insurer-level denial rates for identical services—all patterns we document empirically.

Third, we contribute to a growing literature on fragmentation in the US health care system. Prior work has documented that fragmentation in care delivery—across providers, hospitals, and care settings—increases utilization and costs (Cebul et al., 2008; Agha et al., 2019). We provide evidence that fragmentation on the *insurance* side is a distinct and quantitatively important source of inefficiency. We identify two specific mechanisms through which insurance fragmentation generates administrative frictions. *Fragmented enrollment*—the frequent movement of individuals across insurers due to job changes, life events, and program eligibility transitions (Sommers and Rosenbaum, 2011)—produces sharp spikes in denial rates at the time of switching, as providers and the billing system adjust to new payer rules. *Fragmented rules*—the heterogeneity of coverage policies, bundling conventions, and benefit limits across insurers—produces bimodal denial patterns in which some insurers nearly always deny a service while others rarely do, even for routine procedures. These findings complement research on insurance market structure and consumer frictions (Handel, 2013; Handel and Kolstad, 2015; Shepard and Wagner, 2025) by demonstrating that multi-payer complexity imposes costs not only on consumers choosing among plans but also on the administrative infrastructure through which care is billed and paid.

More broadly, this paper underscores how institutional complexity can undermine the efficiency of market-based health insurance systems. Even in settings with relatively regulated prices and standardized benefits, fragmentation in enrollment and coverage rules can generate substantial administrative waste, uncertainty for patients and providers, and limited gains in cost control. Understanding these frictions is essential for evaluating both incremental reforms and more fundamental changes to the organization of health insurance. By documenting the prevalence and structure of denied claims, this paper aims to provide a foundation for future work on the economics of billing, coordination, and administrative efficiency in health care markets.

2 Background

2.1 Background on Medical Billing and Claims Review

Health insurance billing in the United States involves a multi-stage process that spans patients, providers, and insurers. The process typically begins before care is delivered, when providers verify insurance coverage and eligibility and may obtain prior authorization for selected services. After care is provided, the provider submits a claim to the patient's insurer, listing the services delivered using standardized procedure codes. The insurer then reviews the claim to determine coverage, pricing, and payment, after which it issues payment to the provider

along with an explanation of benefits (EOB) sent to patients. Finally, the provider bills the patient for any remaining balance, including both cost sharing amounts and potentially including bills for denied services. Although this sequence appears linear, in practice it often involves multiple iterations, resubmissions, and appeals, particularly when claims are denied or partially paid.

Claim denials arise during the insurer's review stage, when some or all services on a submitted claim are rejected. Denials most directly affect providers, who must decide whether to absorb the loss, resubmit the claim with additional documentation, or appeal the decision. While denials can also affect patients indirectly (e.g., via receiving an unexpected "balance bill"), the bulk of the evidence suggests that patients bear relatively little cost for denied claims. Even when patients are not financially responsible, they often receive confusing EOB statements indicating that services were not covered, creating uncertainty about whether additional payment will be required. As a result, denials contribute both to administrative burden for providers and to confusion and stress for patients.

High-level industry reports document that claim denials are widespread. The Optum Revenue Cycle Denials Index reports that approximately 12 percent of medical claims nationally were denied in 2023, a rate substantially higher than that for prior authorization, where only about 4–7 percent of claims are subject to review and more than 90 percent of those requests are ultimately approved (Optum, 2024). Reported reasons for denials are diverse. The most common category involves "registration and eligibility" issues, including coordination of benefits across multiple insurers, services delivered to patients who are not covered by the billed plan, or services that exceed benefit limits such as caps on physical therapy visits. Other common reasons include missing or incorrect information on claims, services deemed bundled into broader visits, and services that fall outside an insurer's covered benefits.

2.2 Insurance Fragmentation

A distinctive feature of the US health care system is the fragmentation of insurance coverage across many distinct payers, each with its own administrative rules, provider networks, and benefit designs. Cebul et al. (2008) provide a foundational overview of how organizational fragmentation in US health care—across payers, providers, and care settings—disrupts coordination, degrades care quality, and increases costs. While many goods and services can be efficiently provided through unconnected transactions, health care requires close coordination that fragmented structures undermine.

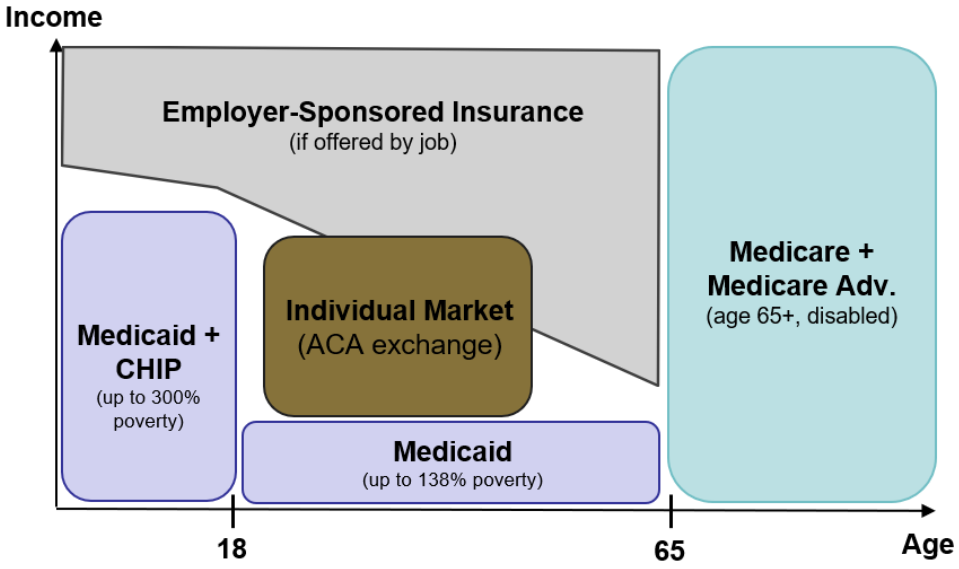
Figure 1a illustrates the fragmented structure of health insurance in the United States. Coverage is organized along two primary dimensions: age and income. Children and low-income adults are covered by Medicaid and the Children's Health Insurance Program (CHIP), with eligibility extending up to 138 percent of the federal poverty level for adults and up to 300 percent

of poverty for children in many states. Working-age adults with access to employer benefits are typically covered by employer-sponsored insurance (ESI), the single largest source of coverage. Those without employer coverage may purchase plans on the individual market through ACA exchanges. Adults over 65 and certain disabled individuals are covered by Medicare, including both Traditional Medicare and Medicare Advantage plans offered by private insurers. These program boundaries create a patchwork of coverage in which a single provider may bill dozens of distinct payers across multiple market segments—each with its own coverage rules, payment rates, network restrictions, and claims processing systems.

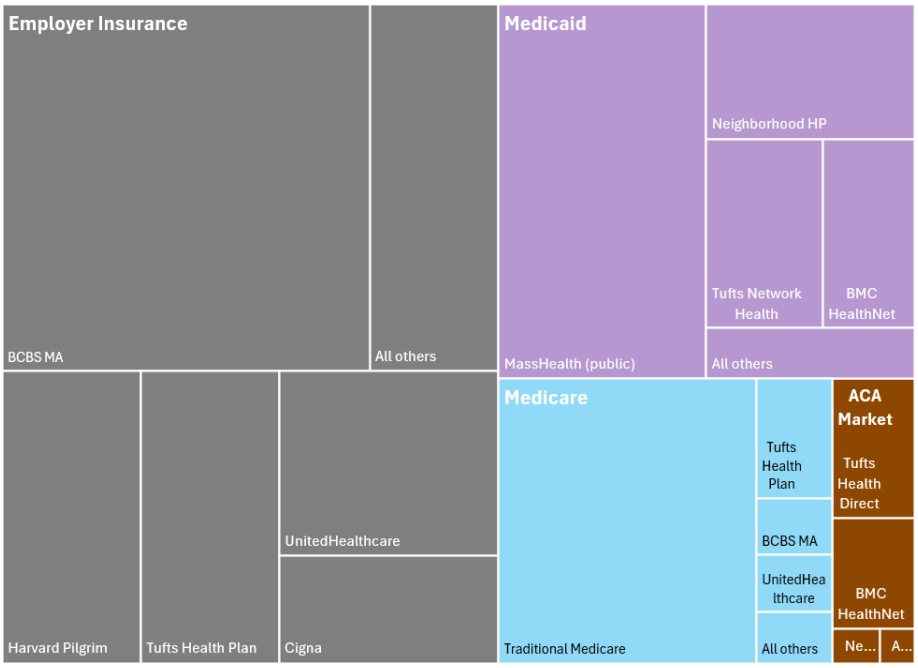
The complexity illustrated in Figure 1a is compounded by the fact that within each market segment, multiple competing insurers operate with distinct administrative systems, as demonstrated in Figure 1b. In Massachusetts—the setting for our empirical analysis—the employer-sponsored market is led by Blue Cross Blue Shield of Massachusetts (approximately 40 percent market share), followed by Harvard Pilgrim Health Care, Tufts Health Plan, UnitedHealthcare, and Cigna, along with numerous smaller carriers (Massachusetts Division of Insurance, 2018). The Medicaid market is split between public MassHealth fee-for-service and several managed care organizations including BMC HealthNet, Tufts Network Health, and Neighborhood Health Plan. The subsidized individual (ACA) market is dominated by Tufts Direct and BMC HealthNet (Massachusetts Health Connector, 2018), while the Medicare Advantage market is led by Tufts Health Plan, Blue Cross Blue Shield of Massachusetts, and UnitedHealthcare (Massachusetts Division of Insurance, 2018). Each of these insurers maintains its own coverage policies, billing rules, provider networks, and claims processing infrastructure, and providers must navigate all of them simultaneously.

This multi-payer, multi-segment structure creates two distinct sources of fragmentation that are central to our analysis. The first is *fragmented enrollment*: individuals frequently move across insurers and market segments due to job changes, income fluctuations, aging into Medicare eligibility, and changes in family structure. These transitions are common—in our data, roughly 10 percent of individuals switch insurers in a given year—and each switch requires both the patient and their providers to adjust to a new set of administrative rules. Sommers and Rosenbaum (2011) estimate that within a year of the ACA's implementation, half of low-income adults would experience a shift in eligibility between Medicaid and exchange coverage, creating administrative disruption for patients and providers alike. The second is *fragmented rules*: even holding enrollment fixed, insurers differ in whether particular services are covered, whether they are paid separately or bundled into broader claims, and whether they are subject to network restrictions or service limits. Providers, who typically bill multiple insurers using the same billing infrastructure, have limited capacity to tailor their claims submission processes to each payer's idiosyncratic requirements.

Figure 1: Fragmented Health Insurance Landscape in Massachusetts



(a) Health Insurance Segments by Age and Income



(b) Massachusetts Health Insurer Market Shares by Segment

Notes: The figure illustrates the fragmented structure of health insurance coverage in the United States and Massachusetts specifically. The top panel shows the main segments of health insurance by age and income group. Each shaded region corresponds to a major insurance program: Medicaid and CHIP (children and low-income adults), the individual market through ACA exchanges, employer-sponsored insurance (ESI), and Medicare (including Medicare Advantage) for adults age 65 and older. The overlapping boundaries between programs highlight the complexity of the coverage landscape that providers must navigate when submitting claims. The bottom panel shows approximate market shares in Massachusetts by insurer (with data as of 2018), grouped by insurance segment.

3 Data

3.1 Massachusetts All-Payer Claims Data

Our primary data are from the Massachusetts All-Payer Claims Database (MA-APCD), which contains medical claims and individually linked enrollment and product characteristics. Our period of analysis covers claims for services delivered from 2013–2020.

Two features of the MA-APCD are central for our purposes. First, the MA-APCD allows us to observe claim denials, which are often not recorded in administrative claims data. Our unit of observation is a charge for a medical service, which may be paid or denied. Because a medical encounter typically consists of multiple distinct services, line item charges for distinct services are administratively bundled into "claims." Following standard practice, we use the term "claim line" to refer to a charge for a specific service, and "claim" to refer to the administrative bundle of services. While claim lines are either paid or denied, a claim consisting of multiple lines may be fully or partially denied. For the MA-APCD, insurers are required to submit all denied lines within partially denied claims. While they are not required to submit fully denied claims, in practice, we observe fully denied claims submitted by most insurers, and we report robustness to various assumptions about the potential exclusion of fully denied claims from our data. In addition to the denial status of a claim line, our data allow us to observe the type of service and the identity of the provider.

Second, the MA-APCD provides broad coverage: the data collect claims submitted by around 60 commercial payers as well as major public programs, including MassHealth (Medicaid) and plans operating in Medicare Advantage. Data submissions are legally required of all insurers licensed by the MA Division of Insurance, including health insurance carriers and their third-party administrators.[1] In addition to Massachusetts residents, these data include all individuals where the site of the insurance contract or product is in Massachusetts.[2] In addition to claims, insurers are required to submit linked individual enrollment and product (plan) characteristics, which allow us to determine product market segments. The data do not include Traditional Medicare claims, and beginning in 2016 (following the Gobeille v. Liberty Mutual Supreme Court decision), do not include all claims from self-insured employer plans.

3.2 Massachusetts Provider Database and Practices

We supplement the claims data with the Massachusetts Health Quality Partners (MHQP) Massachusetts Provider Database (MPD) for 2013–2017. This database provides annual linkages

[1] See state submission instructions: `https://www.chiamass.gov/assets/docs/p/apcd/2026-apcd-submission-guides/2026-apcd-member-simplified-file-submission-guide-FINAL.pdf`.

[2] This covers Massachusetts residents, members of Massachusetts-based employer groups who live out-of-state, and any member of a plan offered by a Massachusetts-based insurer.

between individual clinicians practicing in Massachusetts to practice sites, which we refer to simply as practices; where each site corresponds to a single physical location (address).[3] The MPD allows us to link practices to their claims via physician identifiers (NPIs) listed in the MPD and on claims. A key dimension of practice heterogeneity we analyze is the practice size, which we measure as a number of claim lines submitted by the practice.[4]

4 Descriptive Facts

We begin by establishing three descriptive facts about claim denials. These are (1) denials are prevalent across all market segments, (2) denials are more prevalent for low-priced procedures and are not targeted toward high-cost enrollees, and (3) denials are frequently predictable *ex ante*.

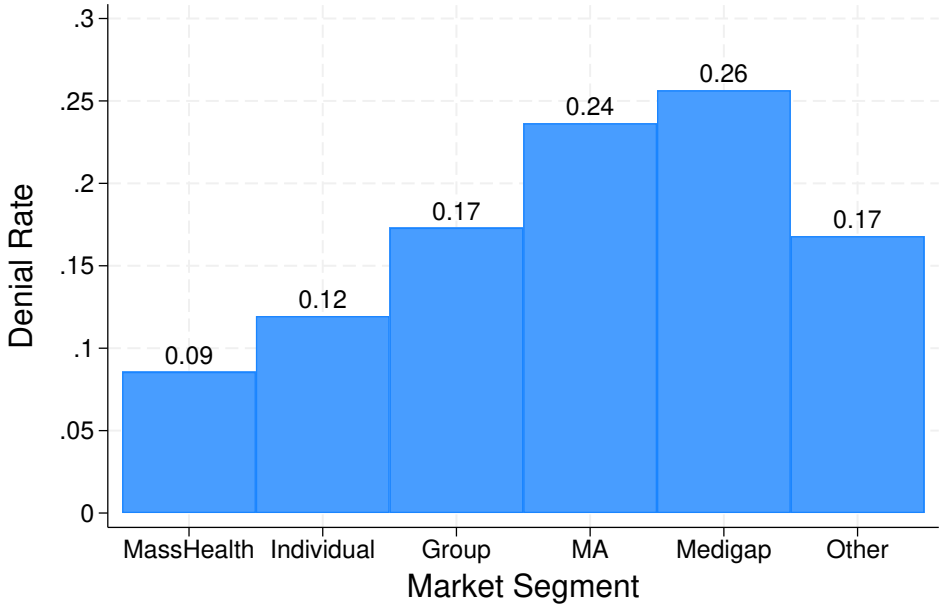
1) Denials are Prevalent First, we show that denials are prevalent. Figure 2 shows the share of claims denied in each market segment. Denial rates are high across all segments, ranging from 9% for MassHealth (Medicaid) to 26% for Medicare Advantage, with an overall rate of 14%. While cross-insurer comparisons should be interpreted with caution—and these figures represent a lower bound on total denials—the key finding is that denials are not confined to public insurance programs but are common across private insurance market segments as well. Appendix Figure A8 shows that from 2013 through 2020, the rate of claim denials changed very little, indicating that denials are a persistent structural feature of the billing system rather than a transitory phenomenon.

2) Denial rates are higher for lower-cost care Second, denials are not generally targeted to high-cost care or high-cost patients. Claim denials are generally framed as a cost-containment tool, with denial of expensive claims often framed as a way of insurers increasing profits at the expense of their enrollees (Alltucker, 2023; Whoriskey, 2024; Miller et al., 2024; Bloche, 2025). As shown by Figure 3a, the probability of a claim being denied is strongly negatively correlated with the price of the procedure. Claims for care in the lowest price ventile are denied at roughly five times the rate of those in the top ventile. This negative relationship holds if we exclude the very cheapest procedures as well, as shown in Appendix Figure A9. Furthermore, this negative

[3]The MPD further describes a hierarchy of practice site affiliations with larger medical groups and even larger "networks" (corresponding to health systems). Because the nature of these affiliations may vary based on the specific contractual arrangement between practices and these other levels of organization, we focus primarily on practices as a unit of analysis.

[4]To construct this measure, we first restrict to claim lines rendered by a physician (NPI) that is a "single-homer" within the focal practice in a given year, to avoid potentially double-counting claims for physicians who work at or switch practices within that year.

Figure 2: Claim Denials by Insurance Market Segment



Notes: The sample is limited to 2016—2020 as these are the only years we have cleaned market segment information. The sample is limited to claims submitted to OrgIDs for which we can identify the insurer and market segment.

relationship is present when looking exclusively at partial or full denials, although it is much stronger for partial denials, as shown by Appendix Figures A4a and A5a.

The negative relationship between price and denial rate does not mean that claim denials are not used to contain costs as providers likely exert more billing effort to avoid claim denials for expensive procedures. At the same time, the relationship does indicate that provider billing effort is much more positively correlated with price than the difficulty of billing for the care. Insofar as claim denials capture the cost of billing for care, this relationship indicates the primary cause of claim denials is unlikely to be providers billing for unnecessary expensive care with additional denials disciplining this behavior.

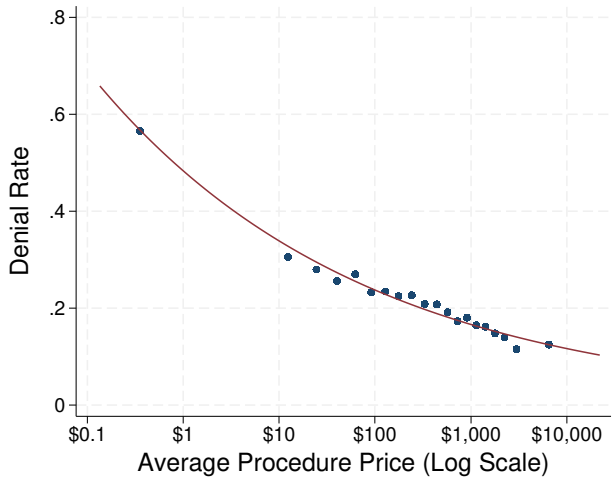
Relatedly, many denials are for common procedures, rather than exotic, expensive procedures. Appendix Table A1 reports the procedures that are most commonly denied. They are routine services. The list is dominated by routine services such as blood draws, basic laboratory panels, and standard evaluation and management codes—precisely the types of high-volume, low-discretion services for which a cost-containment motive is least plausible.

In addition to discouraging high-cost *care*, claim denials may also be used to target high-cost *enrollees*. If experiencing claim denials leads enrollees to change insurers, insurers have an incentive to target claim denials toward unprofitable patients. As shown in Figure 3b, we find no evidence that enrollees who consume more expensive medical care experience claim denials

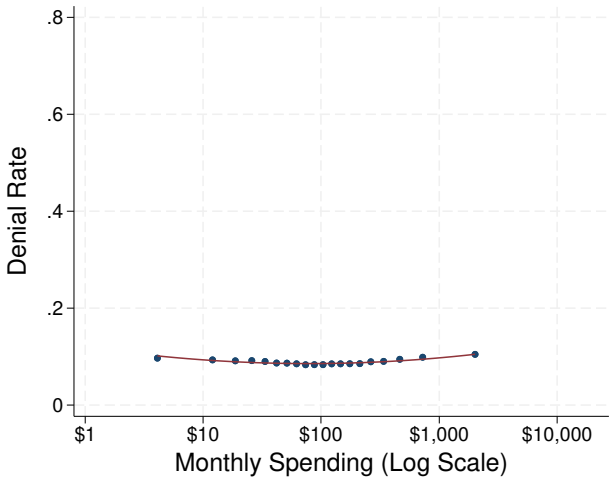
at a higher rate. Higher-spending enrollees do experience more denials in absolute terms simply because they use more health care, but the share of their claims that are denied is no higher than for lower-spending enrollees.

Figure 3: Prevalence of Denials by Cost

(a) Cost of Procedure



(b) Cost of Patient

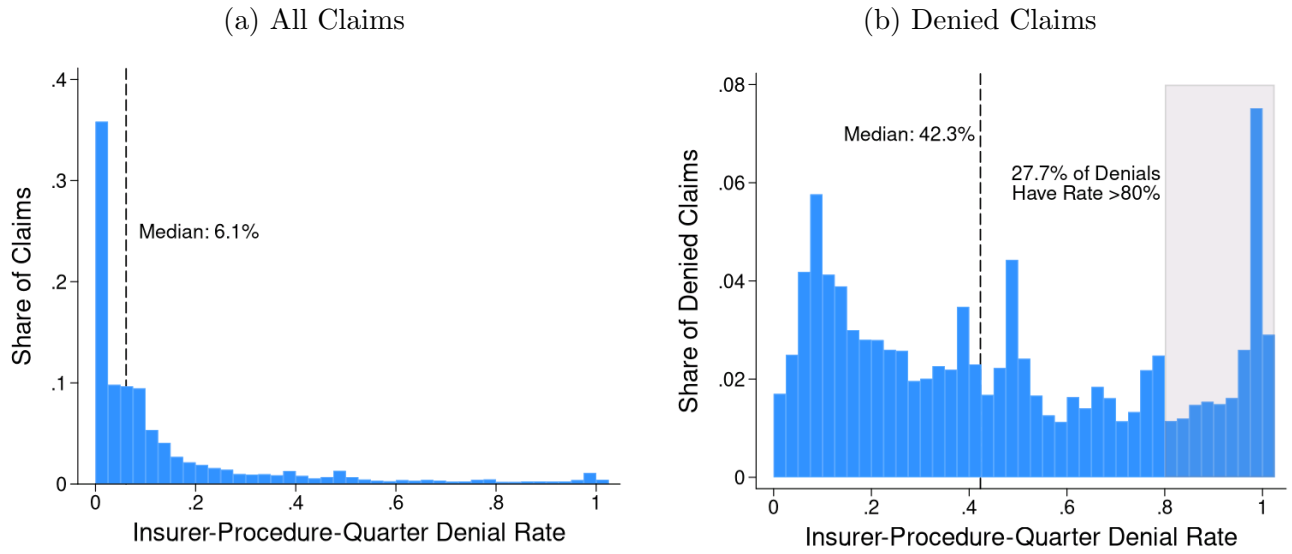


Notes: For panel (a), first, we calculate the price of a procedure by estimating the average paid amount for claim lines for that HCPCS code that are paid a positive amount from 2016–2020. The figure above shows the procedure-level relationship between the share of claim lines denied and the log average price of the procedure, along with a natural-log line of best fit. For panel (b), we limit our data to enrollees in the individual market and to 2016–2020. We calculate total monthly medical spending and the share of claim lines denied over the entire period. The figure presents the average denial rate for each ventile of log spending along with a quadratic line of best fit. Both panels have the same vertical axis scale.

3) Claims for services with near-certain denial probability The third descriptive fact is that denials are often for claims that are nearly certain to be denied *ex ante*. That is, many denials are for procedures that the insurer to which the claim is submitted almost never pays for. Figure 4a shows the insurer-by-procedure denial rate for all claims submitted. We can see that a non-trivial share of claims are submitted to insurers that have very high denial rates for that procedure, with 2.3% of claims being submitted to insurers that deny over 90% of claims for that procedure and 3.3% being submitted to insurers with over an 80% denial rate. That these claims are submitted at all is puzzling: providers should avoid incurring the costs of administering (and billing for) care that is almost certainly not going to be paid for.

Figure 4b highlights that such claims account for a disproportionate share of denials. In this panel we restrict the sample to denied claims, finding that 16.1% of denials are for claims submitted to insurers that deny over 90% of claims for that procedure and 27.7% are for claims the insurer denies over 80% of the time.

Figure 4: Distribution of Insurer-by-Procedure Denial Rates



Notes: Both panels of this figure report histograms of the insurer-procedure-quarter-level denial rate, with each cell weighted by the number of claims in the cell. All cells are limited to those with at least 100 claim lines in the cell. Panel (a) reports the distribution of insurer-procedure-quarter denial rates for all claims from 2013–2020. In panel (b) the sample is further limited to denied claims.

Taken together, these three facts paint a picture of claim denials that is difficult to reconcile with standard explanations. Denials are widespread and persistent, concentrated among low-cost and routine services rather than expensive or discretionary care, not targeted toward high-cost enrollees, and frequently predictable *ex ante*. These patterns raise a natural question: what economic forces can sustain such high levels of seemingly wasteful denials in equilibrium? In the next section, we develop a theoretical framework to address this question.

5 Theory of Claim Denials

In this section, we demonstrate that it is difficult for standard models to explain the descriptive facts about claim denials we presented in the previous section.

The ability for insurers to deny claims is often framed as a way for insurers to control costs (Whoriskey, 2024; Miller et al., 2024; Bloche, 2025; Alltucker, 2023). However, realized (and observed) claim denials are an equilibrium outcome, meaning that regardless of insurer cost control motives, claim denials must be able to persist when providers are responding to the actions of insurers.

Focusing on an auditing model of insurers monitoring whether providers are over-treating patients, we show that with a single provider and single insurer, denials are likely to be rare, cannot be increasing in the financial margin for the provider, and will not occur for procedures

the insurer doesn't cover. We then show that introducing insurer fragmentation along with inability for providers to tailor their strategy to each insurer flips these results and can generate the patterns we see. Finally, we discuss how this model of fragmentation is sustainable in equilibrium.

5.1 Audit Model without Fragmentation

There is an insurer who needs to monitor the claims from a provider. The provider decides whether to administer high treatment, low treatment, or no treatment at all $t \in \{H, L, 0\}$ based on the need of the patient, which can be high or low $\theta \in \{H, L\}$ with high occurring with probability π. The insurer decides whether to audit each claim submitted by the provider $a_t \in \{0, 1\}$ as well as whether to cover each procedure at all $\chi_t \in \{0, 1\}$. The insurer has a style of documentation requirements for the provider to submit the claim, which is a point on the real line ϕ_t, while the provider chooses the style of documentation to provide ψ_t.

Denials occur in three ways. First, if the procedure is not covered, the claim is denied automatically with no auditing. Second, the claim can be denied when the provider is unable to communicate the need of the patient (i.e., does not conform to insurers' documentation requirements) or otherwise has a billing system that is not perfectly compatible with that of the insurer. In this case, the probability of a denial is a function of the distance between the documentation styles of the insurer and provider $d(|\phi_t - \psi_t|)$, with $d(0) = 0$ and $d' > 0$. Third, a claim is denied if it is audited by the insurer and the treatment does not match the need of the patient ($t \neq \theta$). Thus, the probability of denial is given by

$$D(t, \theta, \chi_t, a_t, \phi_t, \psi_t) = \begin{cases} d(|\phi_t - \psi_t|) & \text{if } \chi_t = 1 \text{ and } (a = 0 \text{ or } t = \theta) \\ 1 & \text{if } \chi_t = 0 \text{ or } (a = 1 \text{ and } t \neq \theta) \end{cases}.$$

When a claim is denied, it is unpaid; otherwise, the claim is paid the price of the treatment p_t. The expected payoff for the provider of administering treatment t to patient of type θ is

$$u_P(t, \theta, \chi_t, a_t, \phi_t, \psi_t) = p_t(1 - D(t, \theta, \chi_t, a_t, \phi_t, \psi_t)) - c_t,$$

where c_t is the cost of providing treatment. We assume that high treatment is the one for which the provider's profit margin is greater: $p_H - c_H > p_L - c_L > 0$.[5] The payoff for the insurer is

$$u_I(t, \theta, \chi_t, a_t, \phi_t, \psi_t) = \underbrace{V(t, \theta)}_{\text{Value to insurer}} - \underbrace{c_a \cdot a_t}_{\text{Audit cost}} - \underbrace{p_t(1 - D(t, \theta, \chi_t, a_t, \phi_t, \psi_t))}_{\text{Payment to provider}}.$$

[5] If the profit margin on the low-margin procedure is not greater than zero, the provider will never utilize this procedure.

where c_a is the cost of auditing a claim, and $V(t, \theta)$ gives the insurer's value of the patient receiving the correct treatment:

$$V(t,\theta) = \begin{cases} v_\theta & \text{if } t = \theta \\ v_\theta - l_\theta & \text{if } t \neq \theta \end{cases},$$

where $v_\theta > 0$ and $l_\theta > 0$, meaning that the insurer would prefer patients getting the correct treatment $t = \theta$ (without considering the price of the procedure) over them getting the wrong treatment $t \neq \theta$ and them getting no treatment at all. If $v_\theta \leq 0$, then the insurer would choose not to cover the procedure.

To characterize the equilibria of this game, first note that the provider will never bill for a procedure that is not covered: they incur the cost of treatment c_t and have the claim denied with certainty. Next, note that because the provider's documentation style only affects the probability of denial, the provider will always choose to document in the way that minimizes denials $\psi_t = \phi_t$. Thus, no equilibrium involves denials arising for documentation reasons and the documentation choice by the insurer is strategically irrelevant. Third, there is no equilibrium in which the provider administers appropriate treatment to all patients. This is because the positive margin for the high treatment gives the provider an incentive to overtreat that can only be deterred through the threat of an audit, which if the provider is always treating appropriately cannot be sustained due to audits' costs to the insurer. Finally, there is no equilibrium in which the provider administers low treatment to a high patient if the high treatment is covered,[6] because regardless of the insurer's audit decision, the provider's payoffs are higher for appropriately treating high-type patients. This means that any claims for low treatment will be for low patients, so the insurer will not audit claims for low treatment. This last point implies that, in contrast with the evidence we see in the data, denials should be more common for more expensive procedures.

There are multiple equilibria of this game, but the only one in which denials arise is a semi-separating equilibrium in which the insurer covers the treatment, the provider plays a mixed strategy of overtreating some low-type patients, and the insurer chooses to audit enough claims for high treatment that the provider is indifferent about overtreating.[7] The audit rate for high treatment γ_H and the overtreatment rate α in this equilibrium are

$$\gamma_H = \frac{(p_H - c_H) - (p_L - c_L)}{p_H} \qquad \text{and} \qquad \alpha = \frac{c_a \pi}{(p_H - c_a)(1-\pi)}.$$

In this equilibrium, only claims for high treatment are ever denied, and the denial rate for these

[6] If low treatment is covered and high is not, then there can exist an equilibrium in which the provider "overtreats" high type patients with the low treatment.

[7] This equilibrium is sustainable when $c_a < p_H(1-\pi)$.

claims is

$$\delta_H = \frac{c_a((p_H - c_H) - (p_L - c_L))}{p_H^2}.$$

The key determinant of the audit rate, and thus denial rate, is that the insurer must audit enough claims that the expected payoff to upcoding is equal to the margin for a non-upcoded patient. This means that as the difference in provider profit margins increases, the incentive for the provider to upcode increases, meaning the insurer must audit more frequently and the denial rate increases.

Furthermore, the denial rate even for high treatment claims will generally be quite low. Because the semi-separating equilibrium only arises when $c_a < p_H(1-\pi)$, the denial rate must be less than $(1-\pi)\frac{(p_H-c_H)-(p_L-c_L)}{p_H}$, which will generally be quite low. Making use of the (weaker) condition that the profit margin for the low type procedure must be positive and denoting the markup on this high priced procedure $\mu_H \equiv \frac{p_H-c_H}{c_H}$, we have that the highest possible denial rate is $(1-\pi)\frac{\mu_H}{\mu_H+1}$, which again, will clearly generally be quite low. Thus, in addition to there being no denials for low treatment claims, there should never be 100% denial rates even for high treatment claims.

Thus, this model fails to deliver the patterns we presented in the data in the previous subsection. While we see high denial rates in the data, the model predicts that there will often be no denials at all, and when there are denials, they will generally be low. Second, while we see in the data that denials are decreasing in price, in the model, the opposite is true. Third, while we see realized denial rates of 100% in the data, the model predicts that providers will never bill for non-covered care.

5.2 Audit Model with Fragmentation

Next, we demonstrate that adding insurer fragmentation can flip these results such that the model can match the data. We define **insurer fragmentation** as the two following conditions:

1. The existence of multiple different insurers $j \in \{A, B, \ldots\}$ that differ in coverage rules ($\chi_t^j \neq \chi_t^k$) or billing styles ($\phi_t^j \neq \phi_t^k$).

2. Providers facing a cost to tailoring their billing to each insurer: $h(|\psi_t^j - \psi_t^k|)$.

When both of these conditions hold, we are able to match the patterns in the data.

First, consider the case when the two insurers A and B have the same documentation requirements ($\phi_t^A = \phi_t^B = \phi_t$). In this case, the only equilibria in which the provider administers both high and low treatment and there are realized denials are semi-separating equilibria where the provider and one insurer play mixed strategies and the other insurer plays a pure strategy,

auditing all or none of the claims for high treatment, if covered.[8] Furthermore, if one of the insurers does not cover the low treatment, it may still be worthwhile for the provider to bill for the low treatment in order to get paid by the other insurer for these patients, which can generate a higher denial rate for the low priced treatment.[9] In addition to potentially generating a negative relationship between price and denials, differences in coverage can also result in providers billing for non-covered care, which would generate 100% denial rates at the insurer-by-procedure level, as we see in the data.

Now consider the case where the insurers differ in their documentation requirements $\phi_t^A \neq \phi_t^B$. Here, we are again able to match the patterns we see in the data. In this case, the provider will face the prospect of denials for claims of all levels of treatment and insurers because the provider will choose a documentation style ψ_t between the requirements of the insurers billed. In the simple case where the insurers are identical except for their documentation requirements, the unique equilibrium in which the provider bills for both levels of care is the semi-separating equilibrium described in the previous subsection. The only difference is that now, the denial rate from insurer j for high treatment claims is

$$\delta_H^j = d(|\phi_H^j - \psi_H|) + \frac{c_a((p_H - c_H) - (p_L - c_L))}{p_H^2}(1 - d(|\phi_H^j - \psi_H|))$$

and for low treatment claims is

$$\delta_L^j = d(|\phi_L^j - \psi_L|).$$

The provider will choose ψ_t to minimize these documentation-based denials:

$$\min_{\psi_t} \beta d(|\phi_t^A - \psi_t|) + (1 - \beta) d(|\phi_t^B - \psi_t|),$$

where β is the share of the provider's patients with insurer A. Thus, allowing for differences in documentation requirements can lead to realized denials for even low-priced procedures, particularly if these procedures are ones for which documentation requirements differ the most.

We can further strengthen the ability of this model to generate more denials for low-priced procedures by relaxing the assumption that the provider must bill both insurers using the same style ψ_t and instead allow the provider to tailor the documentation provided at a cost that is increasing in the amount of tailoring $h_t(|\psi_t^A - \psi_t^B|)$. Again assuming identical insurers and, for clarity of the mechanism, assuming each insurer has the same documentation requirements for high and low treatment claims ($\phi_H^j = \phi_L^j$), the equilibrium in which the provider uses both levels

[8] Except in knife-edge cases, the provider can only mix to make one insurer indifferent, meaning there can only be one insurer playing a mixed strategy.

[9] There are two types of equilibria that can potentially generate more denials for the low treatment where both levels of treatment are utilized: one where one insurer covers both levels of treatment and the other only covers high treatment, and one where one insurer covers the high treatment and the other covers the low treatment.

of treatment for enrollees of both insurers, the insurer will tailor the billing strategy more for claims for high treatment because these are claims where denials are the costliest. In particular, the optimal amount of tailoring $\kappa_t \equiv |\psi_t^A - \psi_t^B|$ is given by

$$\frac{\partial u_P}{\partial \kappa_t} = -p_t \left(\beta \frac{\partial \delta_t^A}{\partial \kappa_t} + (1-\beta) \frac{\partial \delta_t^B}{\partial \kappa_t} \right) - \frac{\partial h_t}{\partial \kappa_t} = 0.$$

Because $p_H > p_L$, this means that the optimal κ_H will generally be higher than κ_L, and the amount of tailoring will be greater for the high priced procedure. Under certain conditions, this tailoring effect can outweigh the upcoding effect and there can be higher denials for the low-priced procedure.

To this point, we have discussed how fragmentation can generate the claim denial patterns we see in the data. As further evidence for the importance of fragmentation, we now point out that it is much easier to sustain claim denials in equilibrium with fragmentation. In particular, when prices can adjust to minimize denials, there will be no denials arising from provider upcoding but there can still be denials from the other two sources. In the model, upcoding comes from providers' incentives to overtreat patients for which the lower-margin treatment is appropriate to the higher-margin treatment. Were the financial margins for both procedures the same, there exists an equilibrium with no upcoding and no auditing. The prices can adjust in such a way that both the provider and insurers are better off because audits are costly.

At the same time, because each provider faces a different composition of insurers, their optimal billing styles will differ. Thus, even if an insurer could improve its payoff by changing its coverage rules with one provider, its billing style cannot be changed to avoid denials across *all* providers. In particular, even if each party chooses its billing style ϕ_t or ψ_t optimized in network space with no ability to improve their payoff, unless all insurers and providers can coordinate on the same style, denials will persist. Thus, denials in equilibrium will primarily arise because of a coordination failure.

6 Evidence on Insurer Fragmentation

Having demonstrated theoretically that insurer fragmentation can generate the patterns we see in the data, we next turn to providing positive empirical evidence that fragmentation generates denials. The model yields two testable implications. First, if fragmented enrollment is important, then disruptions to the insurer-provider relationship—such as when enrollees switch insurers—should produce spikes in denial rates as the billing system adjusts to new payer rules. Second, if fragmented rules are important, then cross-insurer variation in denial rates for identical procedures should be large, with bimodal patterns reflecting systematic differences in coverage policies rather than random adjudication noise. Thus, we provide evidence of the importance of

two types of fragmentation: fragmented enrollment and fragmented rules.

6.1 Evidence on Fragmented Enrollment

The fragmentation of the insurance market means that patients often have complex insurance arrangements. As is well-known and studied, patients are commonly uninsured. Less studied are the ways that even the insured are affected by the fragmented insurance system. In particular, individuals regularly change insurers or are insured by multiple insurers. In fact, both of these situations are more common than uninsurance: 10.4% of individuals in our sample switch insurers each year, and more Americans have multiple health insurers than have none (Mykyta et al., 2023).

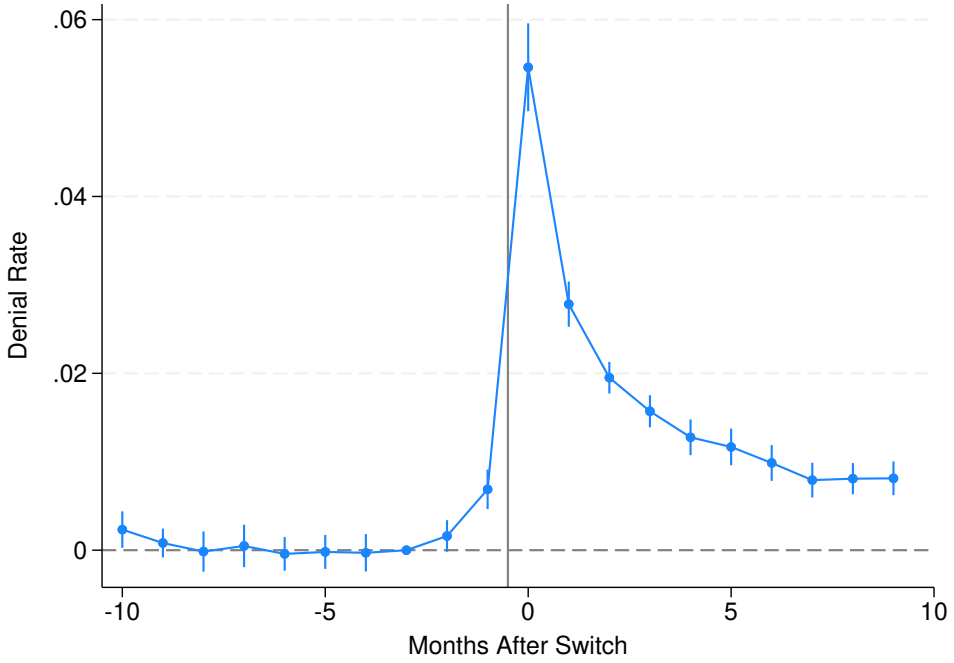
In this section, we present evidence that these complicated, fragmented enrollment situations lead to claim denials. We present two pieces of evidence. First, we show that individuals see a spike in claim denials when they change insurers. Then we show that patients with more than one insurer experience more claim denials than their utilization patterns alone would predict.

Figure 5 presents event-study estimates of the effect of switching insurers on claim denial rates. The figure plots coefficients from a regression of the share of claim lines denied on event-time indicators, individual fixed effects, and insurer-by-month fixed effects, using a sample of "clean" switchers—enrollees who transition directly from one continuous enrollment spell to another with no gaps or overlap. The results show a sharp spike in denial rates at the time of the switch, with denials increasing by approximately 5 percentage points relative to pre-switch levels, a roughly 35% increase. While denial rates begin to decline within a few months of the switch, they remain elevated for a sustained period, suggesting that the adjustment to a new insurer's billing rules is neither immediate nor costless.

Figure A10 shows that most of the acute spike in denials is driven by denials from the outgoing insurer after the patient has changed insurers, while the sustained elevation of denials is entirely driven by denials from the new insurer. In other words, many of the excess denials reflect providers continuing to submit claims to an insurer that no longer covers the patient. This pattern is consistent with lags in the administrative system: providers may not immediately learn that an enrollee's coverage has changed, or they may face delays in updating enrollment information.

Beyond the effects of switching, we also examine whether individuals with more complex concurrent insurance arrangements experience higher denial rates. In the cross section, we see that individuals enrolled in two plans experience a 10 percentage points higher denial rate than those with one insurer. Even conditioning on insurer-by-procedure fixed effects, we find a difference of 6.5 percentage points, or 35% of the baseline. In other words, when two enrollees bill the same insurer for the same care, the enrollee with an additional insurer is substantially more likely to experience a denial. This finding is consistent with coordination-of-benefits complications,

Figure 5: Claim Denials at Insurance Change



Notes: We limit the sample to the 444,513 enrollees continuously enrolled for a spell of at least 10 months with exactly one insurer who immediately transition to another such spell with no gaps or overlap from 2016—2020. This sample excludes most Medicaid patients as many of them are dually enrolled in MassHealth and an MCO. We also keep a control group of enrollees that are part of a continuous enrollment spell of at least 20 months with no switch. This control group is limited to (at most) an equal number of control enrollee-months in each OrgID-by-month cell as there are treated observations in that cell. The coefficients reported in the figure are from a regression of the share of claim lines denied on event-time dummies, individual fixed effects, and OrgID-by-month fixed effects.

in which the presence of multiple payers creates ambiguity about which insurer is primary and increases the likelihood of billing errors.

6.2 Evidence on Fragmented Rules

In this section, we present evidence that fragmented billing rules across insurers lead to claim denials, documenting first that insurers differ in their billing rules in ways that lead to denials before then providing evidence that providers are able to partially learn these rules.

First, we present evidence that insurers differ in their billing rules. To do this, we will first present a few examples of procedures that have widely varying coverage rules to illustrate the types of fragmentation we have in mind, showing that the distribution of denial rates across insurers for these procedures follow a bimodal pattern with some insurers denying nearly all claims for the care while others deny very few. Then, we will document that procedures with this bimodal denial pattern are quite common and comprise a large share of denials, indicating

that insurers vary widely in their coverage and billing rules. Next, we will show that this bimodal pattern is particularly prominent for high-denial procedures, consistent with the theoretical model presented in Section 5.

Finally, we provide evidence that providers' limited capacity to navigate insurer-specific rules can be partially overcome through experience: providers with higher billing volume achieve lower denial rates, suggesting learning-by-doing in the billing process.

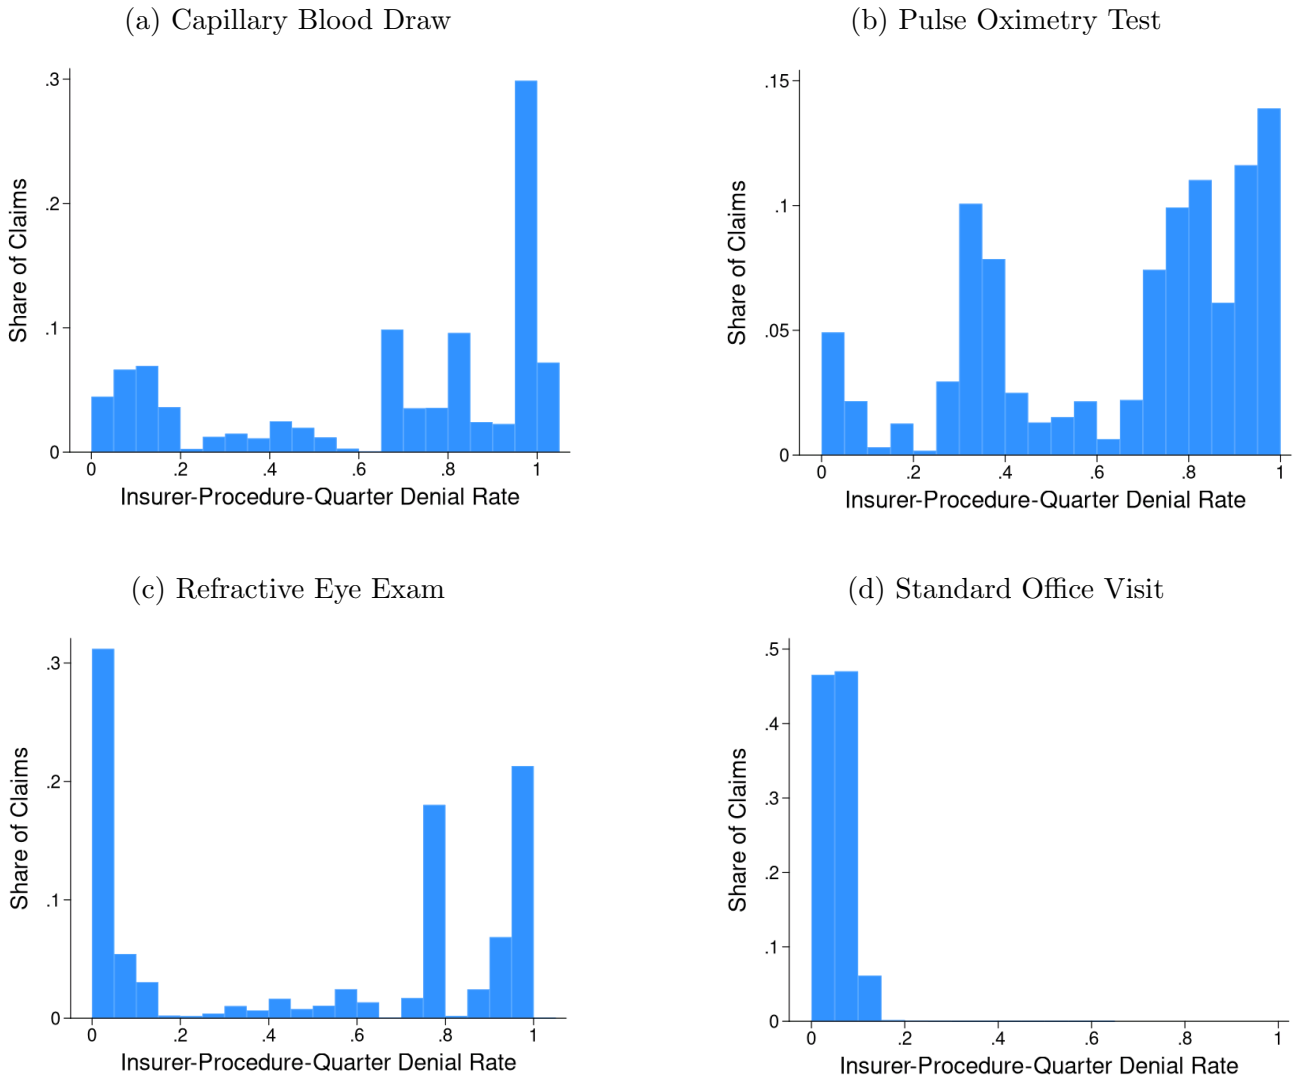
The histograms in Figure 6 below give the claim-weighted denial rates for each insurer for four common procedures: capillary blood draw, pulse oximetry test, refractive eye exam, and physician office visit. The first three of these have a clear bimodal pattern, with some insurers denying the procedure nearly 100% of the time and others having much, much lower denial rates. Each of these illustrates a different common problem stemming from fragmented rules. A capillary blood draw is a pin-prick blood draw for common lab tests, which are often a part of different provider networks. A pulse oximetry test is a common, cheap way to measure the amount of oxygen in a patient's blood. This procedure is often bundled in with the visit itself and not separately paid for. A refractive eye test is the standard way to measure eyes for corrective lenses. This service is sometimes covered as part of a medical benefit and other times is covered by a vision benefit, leading to confusion about which insurer is the appropriate one to bill. Finally, the standard office visit for an established patient is a commonly billed-for code that does not have the same bimodal pattern. Consequently, its average denial rate is less than 5%, while the other three procedures have denial rates above 50%.

In order to test the prevalence of this bimodal pattern, we define "weakly bimodal" and "strongly bimodal" procedures. Both sets of procedures are those whose insurer denial distribution has a Searle's bimodality coefficient of greater than five-ninths, indicating that this distribution is more bimodal than unimodal. We define weakly bimodal procedures to be those where the gap between the 10th and 90th percentile insurer denial rates is at least 10 percentage points. Strongly bimodal procedures are those with a gap of at least 50 percentage points.

We find that bimodal procedures are common and are a disproportionate source of denials. Strongly bimodal procedures represent 21.6% of claims submitted, but 43.1% of denials. Weakly bimodal procedures are 54.5% of claims and 68.6% of denials. The most commonly used bimodal procedures are routine blood collection by venipuncture (36415), therapy procedure using exercise (97110), manual therapy techniques (97140), basic metabolic panel (80048), and emergency department visit (99283). Like the examples above, these suggest different potential sources of fragmented rules: covered separately vs. bundled (venipuncture), provider networks (ED visit), maximum service limits (physical therapy), and service coverage (vision services).

The prevalence of bimodal procedures can explain why providers bill for services that the insurer is extremely unlikely to pay for. In fact, the model predicts that providers will only bill for high-denial rate care when there are other insurers that are much more likely to pay for

Figure 6: Examples of Procedure-Level Insurer Denial Rates



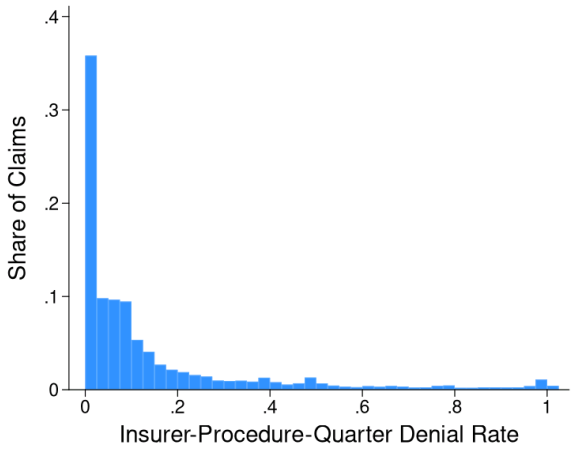
Notes: This figure reports the insurer-level denial rate for four procedures in quarter 1 of 2016. The histograms are weighted by the number of claims submitted to the insurer for that procedure.

it. Figure 7 shows that this is the case: for procedures with higher average denial rates, there remains a meaningful share of insurers that have low denial rates. The figure shows the same insurer denial rate histograms with Figure 7a recreating Figure 4a and showing the distribution for all claims, while Figures 7b, 7c, and 7d limit the sample to procedures with an average denial rate of over 30%, 40%, or 50%, respectively. Even for procedures with an average denial rate above 50%, there are a large number of insurers who rarely deny claims for the procedure.

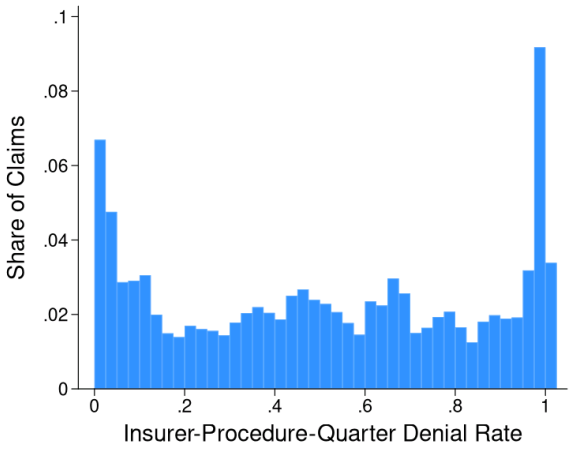
The evidence above establishes that insurers differ substantially in their coverage rules and that providers nonetheless submit claims that are predictably denied. We now turn to evidence that providers can partially learn to navigate these heterogeneous rules through experience.

Figure 7: Distribution of Insurer-by-Procedure Denial Rates by Average Denial Rate

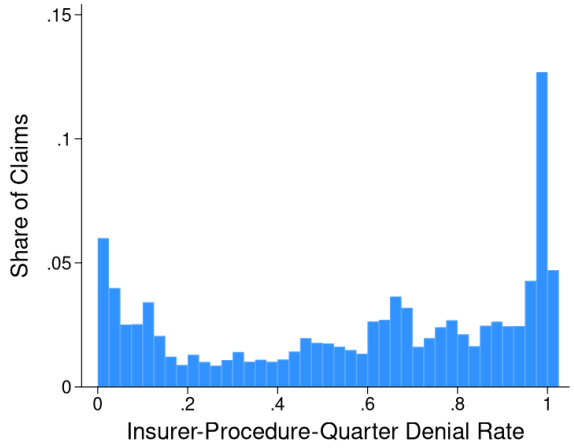
(a) All Procedures



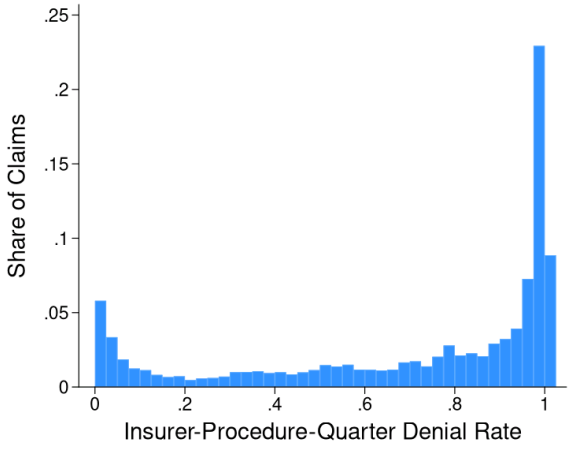
(b) Procedures with a Denial Rate Over 30%



(c) Procedures with a Denial Rate Over 40%



(d) Procedures with a Denial Rate Over 50%

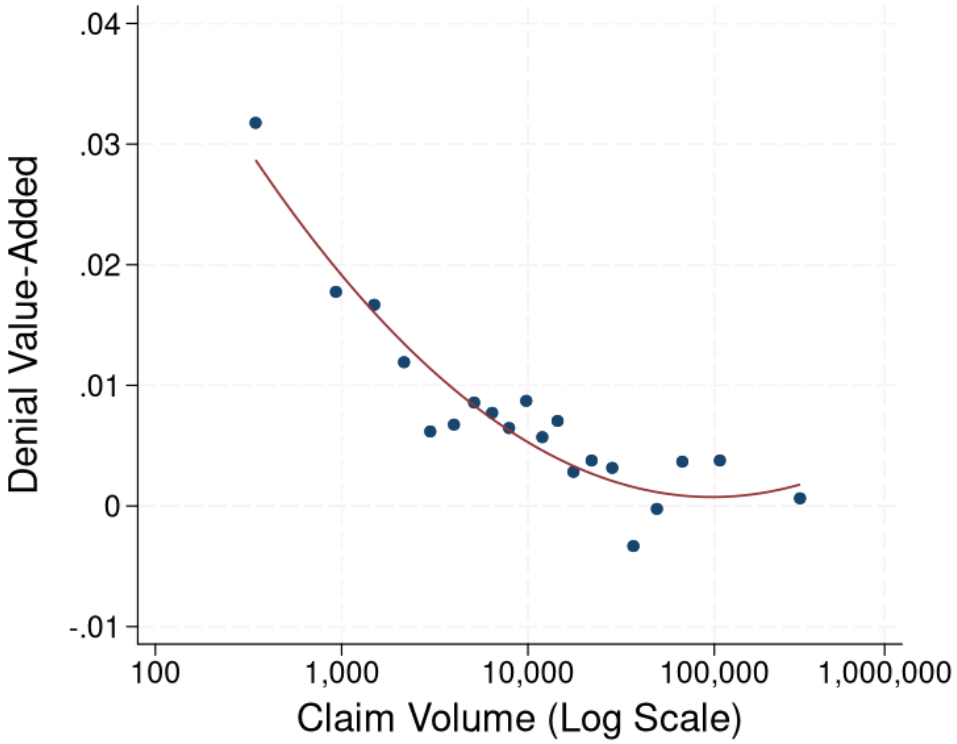


Notes: This figure reports the insurer-procedure-quarter level denial rate for conditional on the procedure-level denial rate in that quarter for 2013–2020. The histograms are weighted by the number of claims submitted to the insurer for that procedure. Panel (a) includes all claims. Panels (b), (c), and (d) are limited to procedures with an overall denial rate of at least 30%, 40%, and 50%, respectively.

Figure 8 shows that providers that bill more face fewer denials, even after adjusting for the mix of payers and procedures for which they bill. This relationship is estimated by computing practice-level value-added in denial avoidance—the residual denial rate after conditioning on procedure-insurer-quarter fixed effects—and plotting it against practice billing volume. The negative relationship is robust: higher-volume practices achieve meaningfully lower denial rates.[10]

[10]We find a similar pattern when estimating value-added at the medical group level. By contrast, the relationship is much weaker when measuring practice/group size by number of physicians or number of affiliated practices, suggesting that it is billing experience specifically—rather than organizational scale per se—that facilitates learning.

Figure 8: Denials by Practice Volume

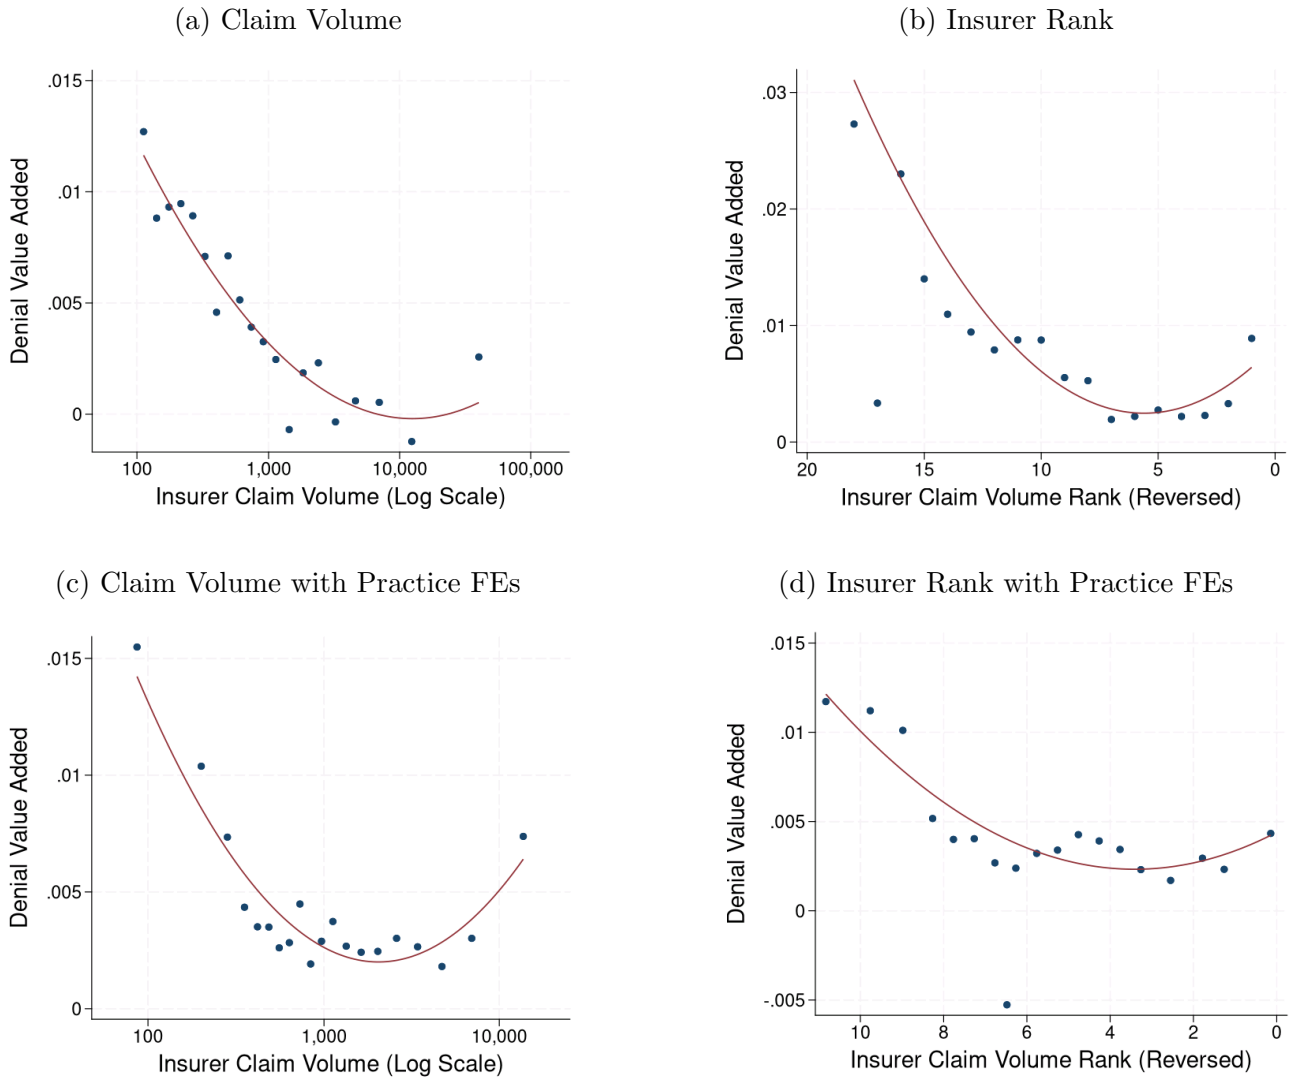


Notes: First, we estimate practice-level value-added by (1) residualizing the share of claim lines denied by HCPCS-orgID-quarter and limiting to cells with at least 100 claim lines, (2) regressing this residualized outcome on practice fixed effects, limiting the data to practice-quarters with at least 100 claim lines. The estimated value of these fixed effects is the practice's value-added, or denial rate above the expected rate conditional on the procedure-insurer-quarter of the claims it is submitting. The figure gives the average denial value-added by ventile of log claim lines submitted by the practice from 2013–2017 along with a quadratic line of best fit.

We also observe a qualitatively similar relationship when measuring experience at the provider-insurer pair level as shown by Figure 9, consistent with the idea that some of the relevant knowledge is insurer-specific. Panel (a) shows that the probability of denial is decreasing (at least up 10,000 claims) in the volume of claims submitted by the practice to the insurer. In panel (b), we see that with the exception of the top-billed insurer, the denial rate is decreasing in the claims volume rank of the insurer within the practice. For example, the denial rate for the second-most through seventh-most billed insurer is over 2 percentage points lower than for the eighteenth-most billed. These patterns are true even when conditioning on practice fixed effects, as shown in Figures 9c and 9d, meaning that within a practice, more claims submitted to a given insurer is generally associated with fewer denials. This suggests that it is billing experience specifically—rather than organizational scale per se—that facilitates learning and fewer denials.

A final piece of evidence for the importance of insurer fragmentation and limited provider ability to navigate a complex insurance system is that vertically integrated providers are more effective at billing their corresponding insurer. Vertical integration has the potential to address

Figure 9: Denials by Practice-Insurer Experience



Notes: First, we use data from 2013–2017 to estimate practice-by-insurer-level value-added by (1) residualizing the share of claim lines denied by HCPCS-orgID-quarter and limiting to cells with at least 100 claim lines, (2) regressing this residualized outcome on practice-by-insurer fixed effects, limiting the data to practice-insurer pairs with at least 100 claim lines. The estimated value of these fixed effects is the value-added, or denial rate above the expected rate conditional on the procedure-insurer-quarter of the claims submitted. The figures gives the average denial value-added by ventile of log claim lines submitted by the practice to the insurer in panels (a) and (c) and the rank of the insurer in terms of claim volume within the practice in panels (b) and (d). In panels (c) and (d), the data is residualized on practice fixed effects. In all panels, a quadratic line of best fit is show in red.

the problems with insurer fragmentation in three ways. First, vertical integration can result in a greater volume of claims exchanged between the insurer and provider, facilitating learning about the insurer's billing rules. Second, vertical integration can dampen the incentives insurers and providers have to capture as much as possible of the potential transfer between them. Third, vertical integration can result in administrative integration, facilitating automated information

exchange.

In line with this potential, we find that vertically integrated health systems have lower denial rates. In particular, we see that vertically integrated providers receive fewer denials from their corresponding insurer than would be predicted given each party's denials when interacting with other insurers and providers. In Massachusetts, there are three large vertically integrated health systems: BMC-HealthNet, Baycare-Health New England, and Partners-Neighborhood, with BMC-HealthNet generally considered the most closely integrated. Regressing an indicator for claim denial on NPI and procedure-by-insurer-by-quarter fixed effects, we find that claims submitted by BMC providers to HealthNet have a 5.1 percentage point (standard error of 0.40) lower denial rate than would be expected given the denials faced by BMC providers when billing other insurers and imposed by HealthNet on other providers. Similarly, the Baycare-Health New England denial rate is 1.9 percentage points (S.E. 0.51) lower than expected. The Partners-Neighborhood denial rate is 2.2 percentage points (S.E. 0.21) higher than expected. Thus, vertical integration is generally associated with fewer denials.

7 Conclusion

This paper documents that denied health insurance claims are a pervasive feature of the US private insurance system and presents evidence that insurance fragmentation—rather than strategic cost containment—is a central driver. Using data from the Massachusetts All-Payer Claims Database, we show that denials are widespread across market segments, concentrated among low-price and routine services, and strongly associated with the heterogeneity of insurer coverage rules. A standard auditing model with a single insurer cannot generate these patterns: it predicts low denial rates, no denials for inexpensive procedures, and no billing for non-covered services. Introducing multiple insurers with heterogeneous rules, combined with providers' limited ability to tailor billing strategies to each payer, fundamentally changes these predictions and can match the empirical patterns we document.

Our empirical evidence points to two specific mechanisms. First, fragmented enrollment—the frequent movement of individuals across insurers—produces sharp spikes in denial rates at the time of switching, driven in part by claims submitted to the wrong insurer but persisting well beyond the transition period. Second, fragmented rules—the heterogeneity of coverage policies, bundling conventions, and benefit limits across insurers—produces bimodal denial patterns for many routine services, with some insurers denying nearly all claims while others rarely deny. These bimodal procedures account for a disproportionate share of all denials. Together, these findings suggest that many denials reflect coordination failures between providers and insurers operating under incompatible administrative systems, rather than the deliberate targeting of high-cost or low-value care.

These results have implications for health care policy. Efforts to reduce administrative waste by standardizing billing rules, harmonizing coverage policies across insurers, or reducing unnecessary plan switching could meaningfully reduce denial rates without restricting access to care. At the same time, our finding that providers with greater billing volume achieve lower denial rates suggests that experience and scale in navigating insurer-specific rules partially mitigate the costs of fragmentation, raising questions about how these costs are distributed across providers of different sizes. More broadly, our analysis highlights that the institutional complexity of multi-payer insurance systems can generate substantial administrative frictions even in the absence of strategic behavior by any individual actor, underscoring the importance of system-level design in determining the efficiency of health care markets.

References

Agha, L., B. Frandsen, and J. B. Rebitzer (2019). Fragmented division of labor and healthcare costs: evidence from moves across regions. *Journal of Public Economics 169*, 144–159.

Alltucker, K. (2023). Is your health insurer using ai to deny you services? lawsuit says errors harmed elders. USA Today.

Becker's Hospital Review (2024). How revenue cycle leaders are using autonomous coding to unlock revenue and reduce denials. *Becker's Hospital Review*. Accessed: 2026-03-10.

Bloche, M. G. (2025). We are all complicit in the big lie of health insurance. STAT News (First Opinion).

Brot-Goldberg, Z. C., S. Burn, T. Layton, and B. Vabson (2023). Rationing medicine through bureaucracy: authorization restrictions in Medicare. Technical report, National Bureau of Economic Research.

Casalino, L. P., S. Nicholson, D. N. Gans, T. Hammons, D. Morra, T. Karber, and W. Levinson (2009). What does it cost physician practices to interact with health insurance plans? *Health Affairs 28*(4), w533–w543.

Cebul, R. D., J. B. Rebitzer, L. J. Taylor, and M. E. Votruba (2008). Organizational fragmentation and care quality in the U.S. healthcare system. *Journal of Economic Perspectives 22*(4), 93–113.

Coase, R. (1960). The problem of social cost. *The Journal of Law and Economics 3*, 1–44.

Crocker, K. J. and J. Morgan (1998). Is honesty the best policy? Curtailing insurance fraud through optimal incentive contracts. *Journal of Political Economy 106*(2), 355–375.

Cutler, D. M. and D. P. Ly (2011). The (paper)work of medicine: understanding international medical costs. *Journal of Economic Perspectives 25*(2), 3–25.

Cutler, D. M., E. Wikler, and P. Basch (2012). Reducing administrative costs and improving the health care system. *New England Journal of Medicine 367*(20), 1875–1878.

Dunn, A., J. D. Gottlieb, A. H. Shapiro, D. J. Sonnenstuhl, and P. Tebaldi (2024). A denial a day keeps the doctor away. *The Quarterly Journal of Economics 139*(1), 187–233.

Eliason, P., R. League, J. Leder-Luis, R. C. McDevitt, and J. W. Roberts (2025). Ambulance taxis: The impact of regulation and litigation on health-care fraud. *Journal of Political Economy 133*(5), 1661–1702.

Forbes Technology Council (2024). Claim denial prediction: Harnessing ai for healthcare revenue cycle management. *Forbes*. Accessed: 2026-03-10.

Gandhi, A. and M. Shi (2025). Screening through soft spending limits: Evidence from the medicare therapy cap. Technical report, National Bureau of Economic Research.

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: when nudging hurts. *American Economic Review 103*(7), 2643–2682.

Handel, B. R. and J. T. Kolstad (2015). Health insurance for "humans": information frictions, plan choice, and consumer welfare. *American Economic Review 105*(8), 2449–2500.

Himmelstein, D. U., M. Jun, R. Busse, K. Chevreul, A. Geissler, P. Jeurissen, S. Thomson, M.-A. Vinet, and S. Woolhandler (2014). A comparison of hospital administrative costs in eight nations: US costs exceed all others by far. *Health Affairs 33*(9), 1586–1594.

Jiwani, A., D. Himmelstein, S. Woolhandler, and J. G. Kahn (2014). Billing and insurance-related administrative costs in United States' health care: synthesis of micro-costing evidence. *BMC Health Services Research 14*(1), 556.

Kaiser Family Foundation (2025, July). Kff health tracking poll: Public finds prior authorization process difficult to manage. `https://www.kff.org/patient-consumer-protections/kff-health-tracking-poll-public-finds-prior-authorization-process-difficult-to-manage`.

League, R. (2024). Administrative burden and consolidation in health care: Evidence from Medicare contractor transitions. *Working Paper*.

Lo, J., M. Long, R. Wallace, M. Salaga, and K. Pestaina (2025, January). Claims denials and appeals in aca marketplace plans in 2023. Issue Brief.

Massachusetts Division of Insurance (2018). Merged market introduction. Technical report, Massachusetts Division of Insurance. Available at `https://www.mass.gov/doc/merged-market-introduction-presentation/download`.

Massachusetts Health Connector (2018). Annual report 2018. Technical report, Massachusetts Health Connector. Available at `https://betterhealthconnector.com/wp-content/uploads/HealthConnectorAnnualReport2018.pdf`.

Miller, T. C., P. Rucker, and D. Armstrong (2024). "not medically necessary": Inside the company helping america's biggest health insurers deny coverage for care. ProPublica.

Mookherjee, D. and I. Png (1989). Optimal auditing, insurance, and redistribution. *Quarterly Journal of Economics 104*(2), 399–415.

Mykyta, L., K. Keisler-Starkey, R. Lindstrom, and L. Bunch (2023, July). People Ages 65 and Older Were Most Likely to Hold More Than One Health Insurance Plan. Section: Government.

Optum (2024). Revenue cycle denials index. Technical report, Optum.

Pozen, A. and D. M. Cutler (2010). Medical spending differences in the united states and canada: the role of prices, procedures, and administrative expenses. *INQUIRY: The Journal of Health Care Organization, Provision, and Financing 47*(2), 124–134.

Shepard, M. and M. Wagner (2025). Do ordeals work for selection markets? evidence from health insurance auto-enrollment. *American Economic Review 115*(3), 772–822.

Shi, M. (2024). Utilization management and administrative frictions. *Working Paper*.

Sommers, B. D. and S. Rosenbaum (2011). Issues in health reform: how changes in eligibility may move millions back and forth between Medicaid and insurance exchanges. *Health Affairs 30*(2), 228–236.

Tseng, P., R. S. Kaplan, B. D. Richman, M. A. Shah, and K. A. Schulman (2018). Administrative costs associated with physician billing and insurance-related activities at an academic health care system. *JAMA 319*(7), 691–697.

Vabson, B., A. L. Hicks, and M. E. Chernew (2025). Medicare advantage denies 17 percent of initial claims; most denials are reversed, but provider payouts dip 7 percent: article examines medicare advantage claim denials. *Health Affairs 44*(6), 702–706.

Whoriskey, P. (2024). Deny and delay: The practices fueling anger at U.S. health insurers. Washington Post.

Woolhandler, S., T. Campbell, and D. U. Himmelstein (2003). Costs of health care administration in the United States and Canada. *New England Journal of Medicine 349*(8), 768–775.

Appendix: For Online Publication

A Full Denials

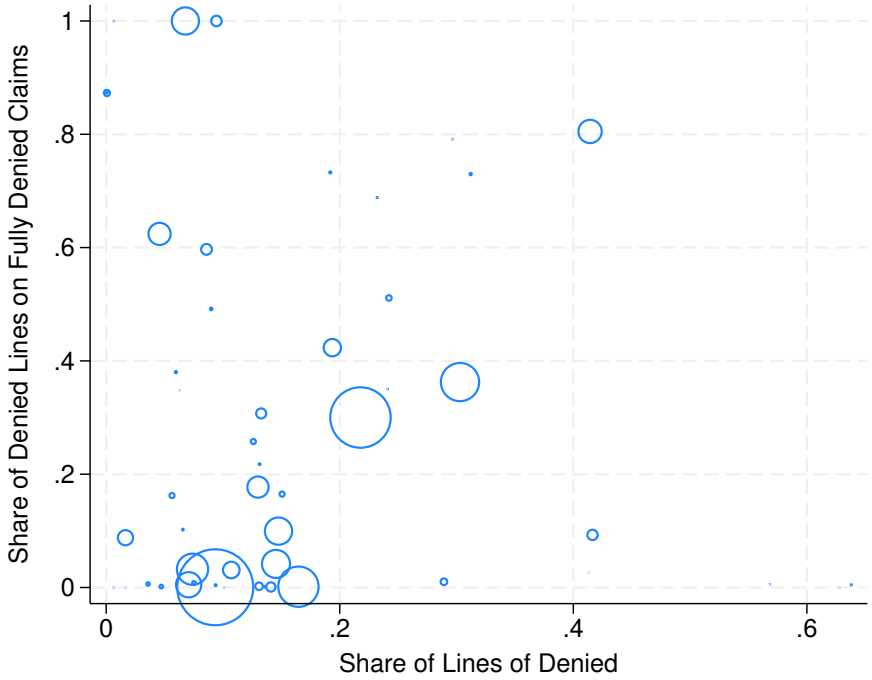
While insurers are required to submit paid and partially denied claims, they are not required to submit fully denied claims. In this appendix, we investigate the extent to which this measurement error matters for our results.

First, we note that insurers differ dramatically in the reported share of denials coming from fully denied claims. Figure A1 reports the share of denials that come from fully denied claims against the overall denial rate in our data. First, we see that despite not being required to report these claims, many insurers do report a non-trivial number of fully denied claims, with some insurers reporting that nearly all of their denials are from fully denied claims.

Second, we see that there is no clear relationship between the share of denials on fully denied claims and the reported denial rate. If insurers only differed in their reporting of fully denied claims (rather than differing in the true underlying denial rate), we would expect a positive, linear relationship between the estimated denial rate and the share of denials from fully denied claims. That is, insurers that report more fully denied claims should appear to have a higher denial rate because they are reporting more claims that go into both the numerator and denominator of the denial rate. We see no such pattern.

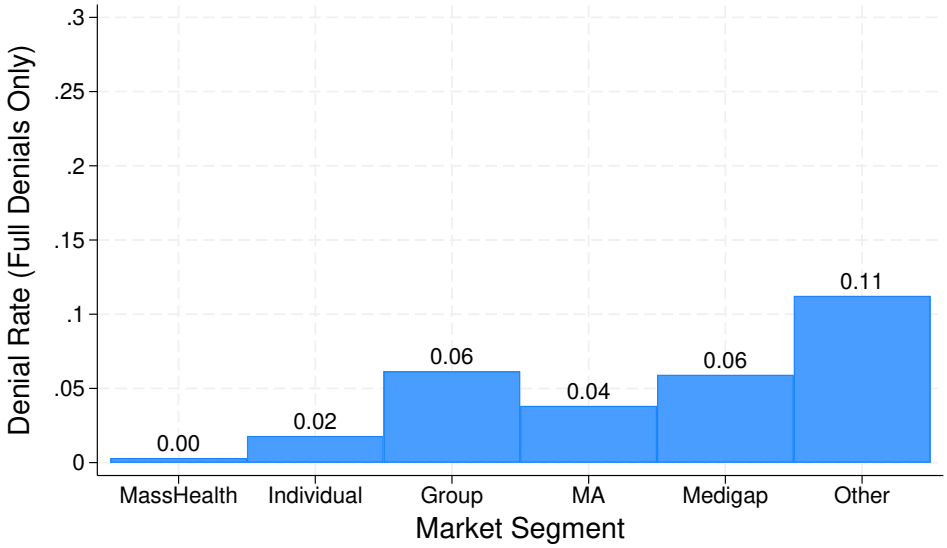
These results indicate that cross-insurer comparisons of the denial rate are somewhat fraught. It is not clear the extent to which differences in the number of fully denied claims stem from true difference in the propensity to fully (rather than partially) deny claims or differences in reporting. For this reason, in the main text, we tentatively treat all of these cross-insurer differences as reflecting true differences in claims processing, while appropriately caveating any cross-insurer comparisons. In this appendix, we break out our results by considering fully denied and partially denied claims separately.

Figure A1: Share of Denials on Fully Denied Claims



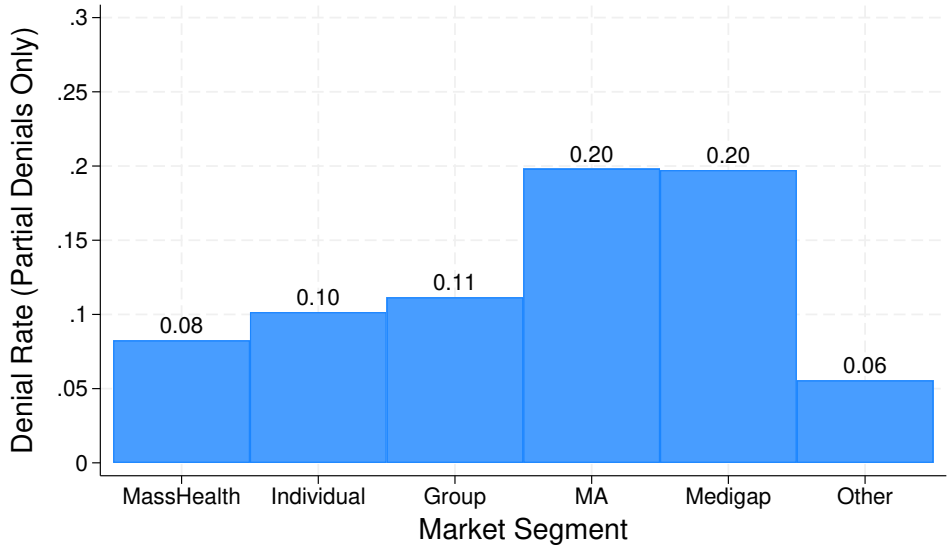
Notes: This figure plots the share of denied lines that are on fully denied claims against the denial rate for each insurer in our data. The size of the circle is proportional to the number of claim lines in the data.

Figure A2: Full Claim Denials by Insurance Market Segment



Notes: The sample is limited to 2016—2020 as these are the only years we have cleaned market segment information. The sample is limited to claims submitted to OrgIDs for which we can identify the insurer and market segment.

Figure A3: Partial Claim Denials by Insurance Market Segment

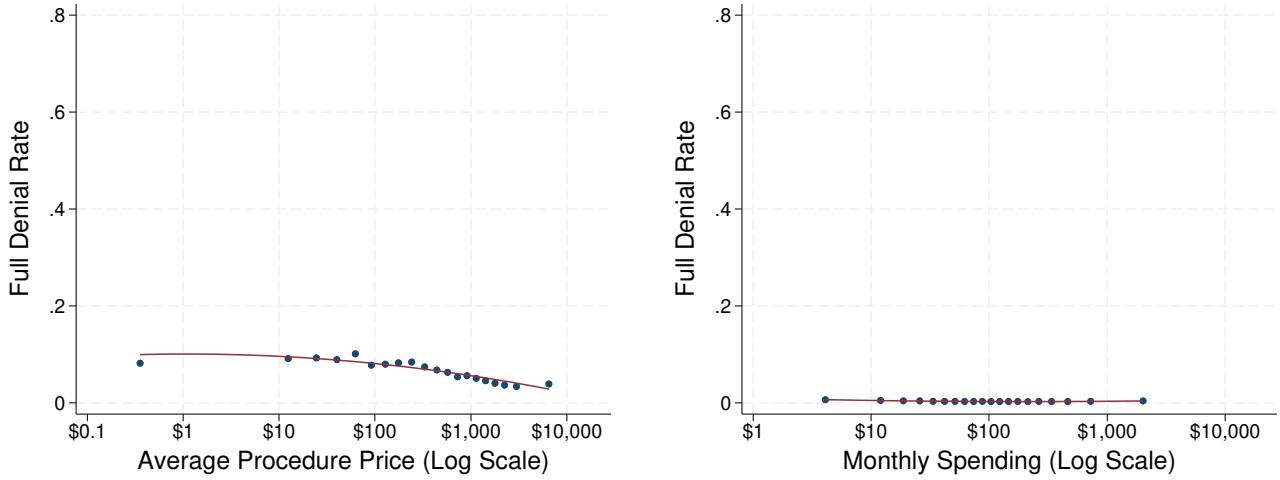


Notes: The sample is limited to 2016—2020 as these are the only years we have cleaned market segment information. The sample is limited to claims submitted to OrgIDs for which we can identify the insurer and market segment.

Figure A4: Prevalence of Full Denials by Cost

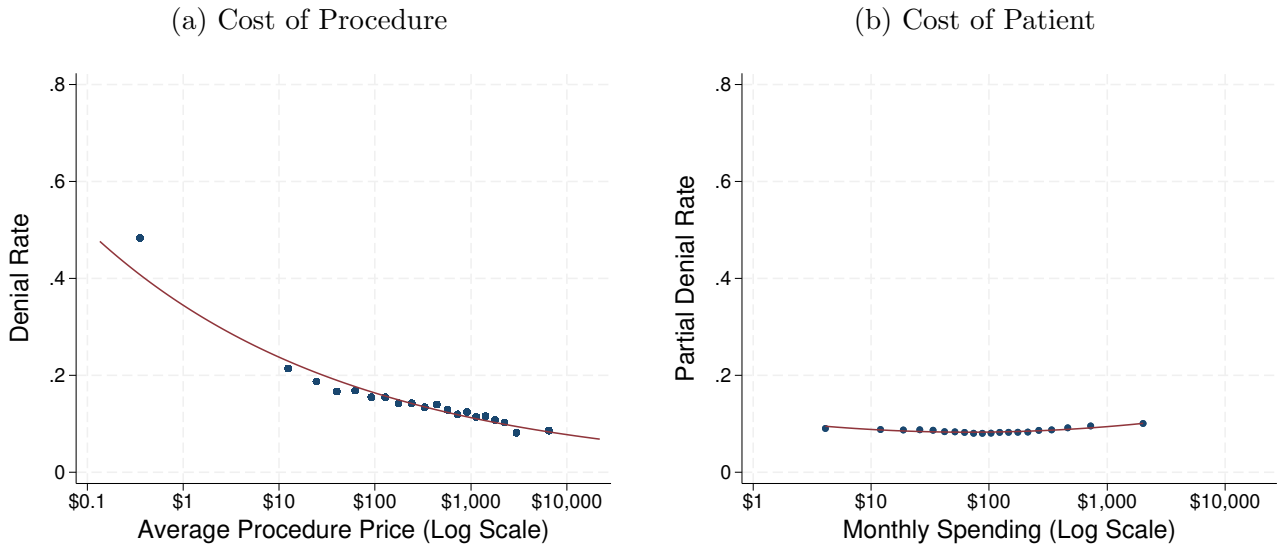
(a) Cost of Procedure

(b) Cost of Patient



Notes: For panel (a), first, we calculate the price of a procedure by estimating the average paid amount for claim lines for that HCPCS code that are paid a positive amount from 2016–2020. The figure above shows the procedure-level relationship between the share of claim lines fully denied and the log average price of the procedure, along with a quadratic line of best fit. For panel (b), we limit our data to enrollees in the individual market and to 2016–2020. We calculate total monthly medical spending and the share of claim lines fully denied over the entire period. The figure presents the average full denial rate for each ventile of log spending along with a quadratic line of best fit. Both panels have the same vertical axis scale.

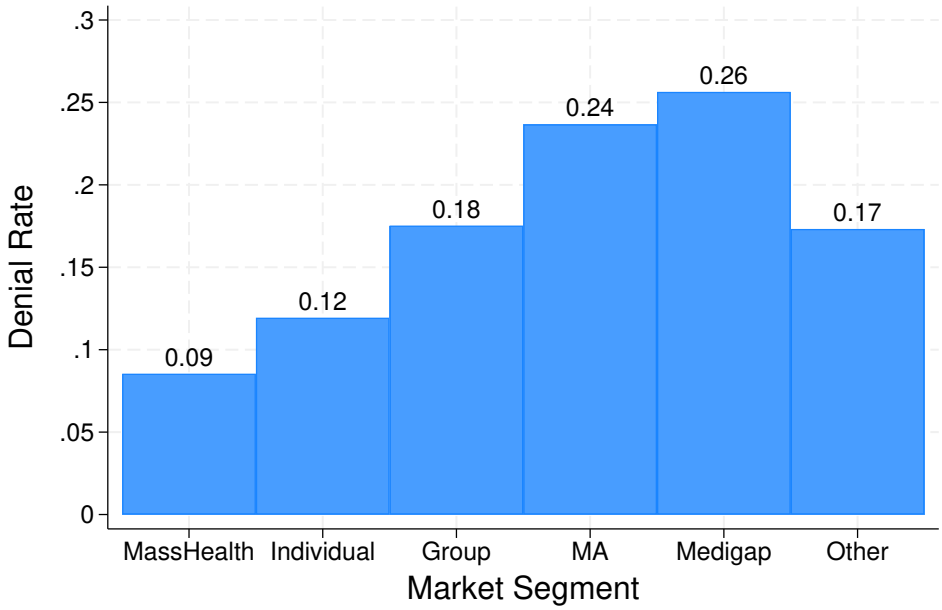
Figure A5: Prevalence of Partial Denials by Cost



Notes: For panel (a), first, we calculate the price of a procedure by estimating the average paid amount for claim lines for that HCPCS code that are paid a positive amount from 2016–2020. The figure above shows the procedure-level relationship between the share of claim lines on partially denied and the log average price of the procedure, along with a natural-log line of best fit. For panel (b), we limit our data to enrollees in the individual market and to 2016–2020. We calculate total monthly medical spending and the share of claim lines partially denied over the entire period. The figure presents the average partial denial rate for each ventile of log spending along with a quadratic line of best fit. Both panels have the same vertical axis scale.

B Excluding Very Small Insurers

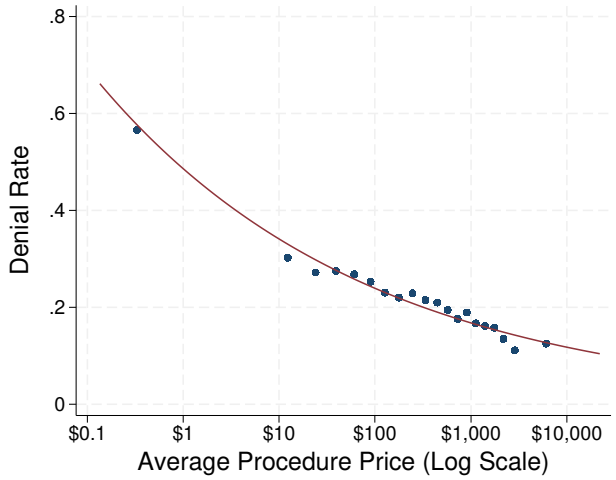
Figure A6: Claim Denials by Insurance Market Segment – Large Insurers Only



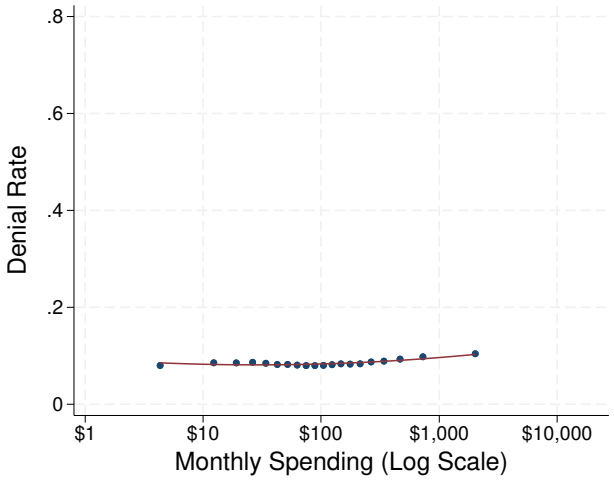
Notes: The sample is limited to 2016—2020 as these are the only years we have cleaned market segment information. The sample is limited to claims submitted to OrgIDs for which we can identify the insurer and market segment.

Figure A7: Prevalence of Denials by Cost – Large Insurers Only

(a) Cost of Procedure



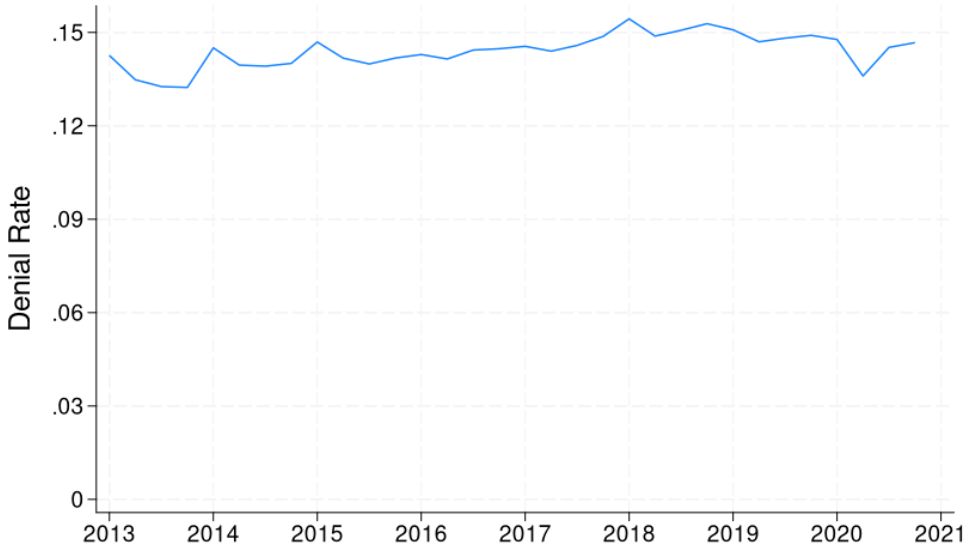
(b) Cost of Patient



Notes: For panel (a), first, we calculate the price of a procedure by estimating the average paid amount for claim lines for that HCPCS code that are paid a positive amount from 2016–2020. The figure above shows the procedure-level relationship between the share of claim lines denied and the log average price of the procedure, along with a quadratic line of best fit. For panel (b), we limit our data to enrollees in the individual market and to 2016–2020. We calculate total monthly medical spending and the share of claim lines denied over the entire period. The figure presents the average denial rate for each ventile of log spending along with a quadratic line of best fit. Both panels have the same vertical axis scale.

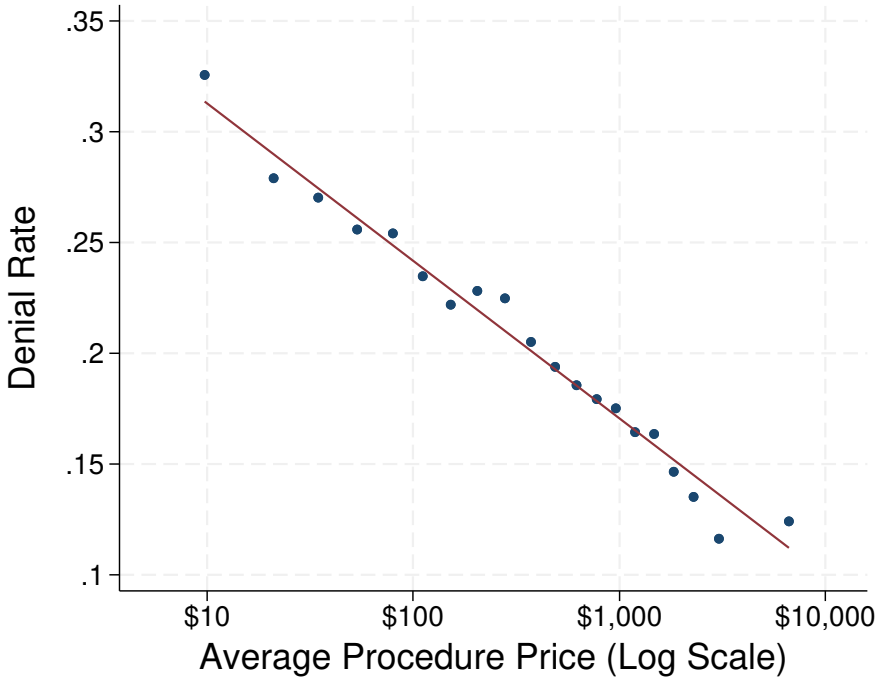
C Additional Results

Figure A8: Claim Denials Over Time



Notes: This figure shows the share of claims denied in the MA-APCD each quarter from 2013 through 2020.

Figure A9: Prevalence of Denials by Cost of Procedure, Excluding Very Cheap Procedures



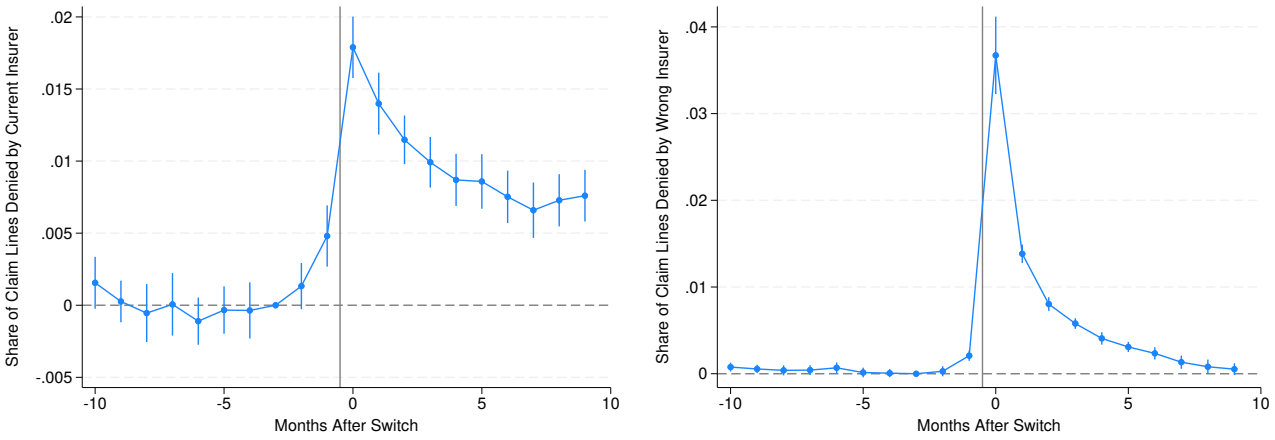
Notes: This figure recreates Figure 3a excluding procedures with a price less than $5 and using a linear line of best fit.

Table A1: Most Commonly Denied Codes

Code	Type	Description	Denied Lines	Denial Rate	Mean Price
36415	Lab test	Collection of venous blood by venipuncture	19,048,827	0.48	$8.24
85025	Lab test	Complete blood count (CBC), automated w/ differential	4,937,538	0.21	$16.89
99213	Eval & Mgmt	Office visit for established patient (20–29 min)	4,144,662	0.05	$76.77
92015	Ophthalmology	Determination of refractive state	4,040,844	0.52	$153.17
T1019	Home health	Personal care services (per 15 mins) – Non-medical asst. with ADLs	3,914,315	0.04	$148.47
97110	Therapy/Rehab	Therapeutic procedure, one or more areas, 15 min	3,282,375	0.09	$42.04
97140	Therapy/Rehab	Manual therapy techniques, 15 min	3,277,502	0.15	$28.32
80053	Lab test	Comprehensive metabolic panel	2,950,996	0.19	$39.97
90686	Immunization	Influenza vaccine (intramuscular injection)	2,859,373	0.33	$22.55
99214	Eval & Mgmt	Office visit for established patient (30–39 min)	2,825,735	0.04	$105.11

Notes: This table reports the ten most commonly denied procedures in quarter 1 of 2016 along with the number of denied lines, the share of lines denied, and the average price of the procedure conditional on the claim being paid a positive amount.

Figure A10: Claim Denials at Insurance Change by Source



(a) Denials from Correct Insurer

(b) Denials from Wrong Insurer

Notes: We limit the sample to the 444,513 enrollees continuously enrolled for a spell of at least 10 months with exactly one insurer who immediately transition to another such spell with no gaps or overlap from 2016—2020. This sample excludes most Medicaid patients as many of them are dually enrolled in MassHealth and an MCO. We also keep a control group of enrollees that are part of a continuous enrollment spell of at least 20 months with no switch. This control group is limited to (at most) an equal number of control enrollee-months in each OrgID-by-month cell as there are treated observations in that cell. The coefficients reported in the figure are from a regression of the share of claim lines denied on event-time dummies, individual fixed effects, and OrgID-by-month fixed effects, with the denials in the first panel limited to those received from the insurer with which the enrollee was currently enrolled while the second panel considers denials from all other insurers.