Arbitration with Regulatory Capture: Evidence from Disputes over Surprise Medical Bills

Benjamin Chartock* *Bentley University*
Christopher Whaley† *Brown University*

September 11, 2026

DRAFT: NOT FOR CITATION OR CIRCULATION

Abstract

Frequent surprise out-of-network medical bills prompted a 2020 law establishing what is likely largest final-offer arbitration system in existence. Unlike canonical final-offer arbitration, outcomes here are lopsided. Providers win 84 percent of disputes, prevailing offers average three times typical in-network rates for the same services, and case volumes vastly exceed projections. We develop a model in which arbitrators, compensated per case, face incentives that tilt awards toward providers. Such incentives alone would relocate the equilibrium upward while preserving balanced win rates. Asymmetric exposure breaks that balance: because an insurer that raises its arbitration offer risks conceding higher rates across its entire network, it bids below the arbitrator-favored level even at the cost of losing. Our model thus accounts jointly for inflated awards and one-sided outcomes, neither of which either mechanism delivers alone. Using a difference-in-differences design on claims data from a large national health insurer, we find negotiated in-network rates increase by 27% for providers exposed to the No Surprises Act relative to unexposed providers, demonstrating that the law's impact extends beyond arbitrated claims to health care costs across the entire privately insured market.

*bchartock@bentley.edu, †christopher_whaley@brown.edu. We thank Zarek Brot, Michael Chernew, Stuart Craig, Leemore Dafny, Jacob Dorn, Erin Duffy, Erin Fuse Brown, Tal Gross, Jon Gruber, Anup Malani, Lindsay Murtagh, Steve Parente, Seth Richards-Shubik, Kosali Simon, and Philip Valtadoros, as well as seminar participants at Johns Hopkins Department of Economics, ZEW/Mannheim workshop, Indiana University, 2025 ASSA Annual Meeting, APPAM, Health and Human Services (ASPE), Bates White, and the University of Wisconsin Department of Risk & Insurance for helpful comments. Neil Mehta, Nandita Radhakrishnan, Qi Lei, and Chris Liu provided excellent research assistance. All remaining errors are our own. Funding provided by Arnold Ventures and Patient Rights Advocate.

1 Introduction

Arbitration is used to resolve disputes where contracts are incomplete ex ante (Williamson, 1975; Grossman and Hart, 1986), which is why it governs settings as varied as union wages, consumer contracts, and international trade. Its appeal rests on a strong theoretical foundation: final-offer arbitration is self-disciplining, because each party is punished for asking too much, so offers converge and neither side wins systematically (Farber, 1980; Ashenfelter and Bloom, 1984). We study what is likely the largest final-offer arbitration system ever instituted, established by the No Surprises Act, and find the opposite. Providers win 84 percent of disputes, and prevailing offers run roughly three times the statutory benchmark. This is not arbitration failing at random but more like a mechanism design failure with a testable signature: arbitrators compensated per case have reason to raise where the equilibrium sits, while insurers, protecting the in-network rates they negotiate elsewhere, decline to follow them upward. High awards and lopsided wins are the joint product, and neither mechanism alone delivers both.

A patient who goes to an in-network emergency department may be treated by an out-of-network emergency physician, anesthesiologist, or radiologist they never selected, and receive a bill for the difference. Because the patient cannot shop at the point of care, some rate must be settled after the fact where none was agreed to in advance. Before the law, the clinician could simply bill the patient directly for the balance (holdup, in a specific market). And this is not rare: roughly one in five emergency visits involves an out-of-network clinician (Garmon and Chartock, 2017; Cooper, Scott Morton, and Shekita, 2020). Congress, drawing on a number of state laws, held patients harmless from these bills and created a dispute resolution mechanism relying on final-offer arbitration to settle payment between insurer and provider. Because health care prices in the private insurance market are bilaterally negotiated, the arbitration mechanism defines the disagreement payoff in a Nash bargaining sense. Raising the arbitration payoff therefore raises the price at which a provider is willing to stay in network, so the law moves rates on care that never enters a dispute, as well. The Congressional Budget Office projected roughly twenty thousand

disputes per year; more than three million have been filed in the system's first three years, at an administrative cost of about five billion dollars (Hoadley and Watts, 2025).

The system that produced these disputes also produced an unusual amount of evidence about them. CMS publishes each dispute's outcome together with both parties' final offers and the names of the insurer, the provider, and the arbitrator who decided it. Every offer is denominated against the qualifying payment amount, so bids are comparable across cases, services, and markets, a normalization the older data never had. That matters because the canonical final-offer literature rests on New Jersey police and fire disputes from the 1970s, 1980s, and 1990s (Farber, 1980; Ashenfelter and Bloom, 1984), and even recent theoretical work (Kong, Silveira, and Tang, 2025) continues to run on that same data, so the tests that would distinguish competing theories of arbitrator behavior have never had a setting to run in. Publication is not merely a convenience for the researcher. Because offers appear with names attached and at volume, a party's bid is informative to counterparties outside the dispute in which it was made, which is precisely why an insurer's arbitration offer carries consequences across its entire network. The feature that makes this system observable is the same feature that distorts it. We observe more than half a million disputes with both sides' bids.

We begin with the canonical Farber framework and rule out two potential explanations for deviation from this model in the data. Risk aversion, which can cause deviations from 50% win rates, cannot be rationalized by any parameters found in the literature. And a provider-friendly arbitrator alone does not generate uneven win rates, because both parties re-center their offers around the new standard. Accordingly, the win rate is a puzzle that survives even inducement by the arbitrator toward higher payment amounts.

The resolution of this puzzle lies in the fact that publication of insurer and provider offers represents a verifiable signal of willingness to pay or be paid for a given service in a given market. An insurer's bid can be used by other providers to exert leverage on in-network negotiations, whereas the provider groups have little in-network business at stake, thus their bids bind less. The insurer therefore bids low and loses often, as losing cases can be a cheaper alternative than re-setting network rates to higher levels. Arbitrator incentives

then set where the equilibrium sits: arbitrators compensated on a fee-for-service basis (per case payment) compete for the caseload that providers originate, which tilts the standard upward, with or without any ownership stake in the arbitrator by a firm that also owns a stake in providers bringing cases. Together, the two mechanisms deliver what neither can alone; exposure asymmetry to the remaining book of business produces lopsided win rates, while arbitrator incentives produce inflated awards. Where exposure is symmetric, e.g., a large provider system against a dominant insurer, balance should return.

Section 4 takes these predictions to the public CMS data. Using the final offers and who was selected from each dispute, we estimate the arbitrator's ideal outcome directly, and find it sits well above the benchmark rate; provider offers cluster near three times the QPA. Insurer offers cluster near the benchmark itself, and awards tend to go to the provider, producing the 84% win rate. The outcomes are not uniform and vary systematically across arbitration entities, the pattern competition for caseload predicts. Insurer behavior matches the exposure mechanism, as they submit an offer exactly at the benchmark in more than a third of disputes, including cases they expect to lose, the protest behavior our model generates when losing a dispute is cheaper to the insurer than conceding a signal to all in-network providers. And where exposure is symmetric, balance returns. Disputes between large hospital systems and dominant insurers run close to even, exactly where the model says the asymmetry should vanish.

The model's final prediction is that the arbitrator's preferences for ideal outcomes subsequently pass through to negotiated prices, affecting the disagreement payoff in Nash bargaining in the way laid out by Prager and Tilipman, 2025. Because arbitration defines that payoff, a provider's improved arbitration outlook should raise the rates it commands in-network, including for care that never goes to arbitration. Section 5 tests this using data from a large national health insurer in a difference-in-differences design, comparing in-network rates for providers exposed to No Surprises Act arbitration against unexposed providers, before and after the law. In-network rates for providers exposed to the NSA rise by 27% relative to unexposed providers. The law's effects therefore are not confined to the small fraction of claims that reach arbitration; they appear in negotiated prices across the

insurance market.

The disputes themselves are concentrated among a small set of repeat players. A handful of private equity backed physician staffing firms account for roughly half of all filed cases, and private equity firms hold ownership stakes in at least five of the approximately fifteen certified arbitration entities. This is of interest in part because a growing literature (Atul Gupta, Howell, Yannelis, and Abhinav Gupta, 2024), (Atul Gupta, La Forgia, and Sacarny, 2024), (Singh, 2022), (Singh, Radhakrishnan, Adler, and C. Whaley, 2025) examines whether private equity ownership changes firm behavior in health care. Economically, the ownership form is less important than the behavior. Private equity is known to have economic benefits, yet real or perceived conflicts may erode some of that benefit. In at least one instance, the same firm holds stakes in both a provider filing cases and an arbitration entity deciding them. Our model does not require an ownership link as competition for cases tilts any fee-for-service arbitrator towards parties who originate cases. Common ownership is simply the mechanism's sharpest expression. Nor is the mechanism confined to health care. Theoretically, the distortion we identify comes from named parties and posted offers at high volume, which reprice business outside the dispute. Mandatory arbitration now is found in most employment agreements[1] and contracts attached to consumer goods like cell phones and credit cards, so the findings here are relevant beyond health care and across domains where arbitration displaces courts in resolving incomplete contracts.[2]

The remainder of the paper proceeds as follows. Section II develops the model of final-offer arbitration with published offers and derives its predictions for win rates, award levels, and pass-through. Section III describes the institutional setting and the data. Section IV presents the arbitration results. Section V presents the effects on negotiated in-network prices. Section 6 concludes.

[1] For example, 56.2% of non-union private sector employees are subject to mandatory arbitration in the US, a large increase from the approximately 2% of workers that were subject to this practice in the early 1990s.https://files.epi.org/pdf/144131.pdf

[2] https://files.consumerfinance.gov/f/201503_cfpb_arbitration-study-report-to-congress-2015.pdf

2 Model

This section develops a model of No Surprises Act final-offer arbitration. We begin with the standard benchmark, which delivers symmetric equilibrium outcomes and equal win rates. We then establish a neutrality result that is central to what follows: within the standard framework, *no* shift in the arbitrator's preferences can generate asymmetric win rates, and asymmetric risk aversion can do so only under implausible parameter values. Breaking the symmetry requires breaking the symmetry of the parties' *stakes*, not the location of the arbitrator's ideal point.

We then introduce the mechanism that does so. Under the NSA, the Centers for Medicare and Medicaid Services publishes each party's final offer, by name, for every dispute. An offer is therefore not a free action: it is a public statement of what a firm is willing to pay or accept, observed by every counterparty that firm negotiates with outside the dispute. Because insurers hold enormous in-network books relative to their arbitration volume, and because the provider groups that dominate NSA arbitration hold almost none, this consideration binds far more tightly on insurers. The resulting equilibrium generates lopsided win rates, a point mass of insurer offers at exactly the qualifying payment amount, and a prediction about negotiated in-network prices that we take to the data.

2.1 Baseline Model of Final-Offer Arbitration

We begin by introducing the game of final offer arbitration that aligns with the standard model in the literature (e.g., Farber, 1980; Ashenfelter and Bloom, 1984). This final-offer arbitration game has three actors: an insurer, a provider, and an arbitrator, indexed by i, p, and a, respectively. We use the notation w (wage) to refer to the reimbursement paid from the insurer to the provider. In the model, insurers and providers offer w_i and w_p, respectively, to the arbitrator. Throughout we normalize the qualifying payment amount to one, so all offers and awards are expressed as multiples of the QPA.

2.1.1 Arbitrator's Utility Function & Timing

The timing of the game is as follows. First, the insurer and provider simultaneously offer w_i and w_p to the arbitrator, who then chooses one of the two offers based on his preferred outcome w_a. The arbitrator's utility function is modeled as a quadratic loss from the distance between his preferred outcome and what is eventually chosen, $U_a = -(w_a - w)^2$, which leads to a simple decision rule where the arbitrator picks the offer closer to w_a. The arbitrator selects the insurer's offer if w_a is below the midpoint of the two offers.

2.1.2 Information Structure

We follow the literature and assume that the arbitrator knows with certainty his preferred outcome w_a conditional on the observable facts of the case, but that the parties only know the arbitrator's preferences up to a distribution. Parties understand this uncertainty as a random variable with cumulative distribution function $F(\cdot)$ and probability density function $f(\cdot)$. Writing

$$m \equiv \frac{w_i + w_p}{2}, \tag{1}$$

the probability that the insurer's offer is chosen is

$$\Pr(\text{insurer wins}) = \Pr(w_a \leq m) = F(m), \tag{2}$$

and the provider's offer is chosen with probability $1 - F(m)$.

We emphasize one feature of the information structure that we will maintain throughout and that distinguishes our explanation from others in the literature: F is common knowledge. Both insurers and provider groups in this setting are sophisticated repeat players who observe published arbitration outcomes quarterly, and nothing in what follows requires either side to misjudge the arbitrator.

2.1.3 Strategies & Equilibrium

Naturally, insurers want to minimize their payment, conditional on the provider making his optimal offer, and the provider wants to maximize this payment, again conditional on the insurer's best response. A pair of offers (w_i^*, w_p^*) is a Nash equilibrium if w_i^* and w_p^* solve, respectively,

$$\min_{w_i} \quad w_i\, F(m) + w_p\, [1 - F(m)] \tag{3}$$

$$\max_{w_p} \quad w_i\, F(m) + w_p\, [1 - F(m)] \tag{4}$$

Solving via first-order conditions:

$$F(m^*) = \frac{1}{2} \tag{5}$$

$$w_p^* - w_i^* = \frac{1}{f(m^*)} \tag{6}$$

In equilibrium, the median of the arbitrator's preferred distribution is at the average of the offers, and the offer gap gets bigger as the distribution gets more spread out. The incentives in this game are thus quite similar to a sealed-bid, first-price auction: a lower bid saves the bidder money but reduces the probability of the bid winning. A provider trades off offering slightly higher, and earning more if he wins, against reducing the probability that his offer is selected. The insurer has a symmetric and opposite incentive to bid low, but not so low that the arbitrator thinks he is being unreasonable and the other party prevails.

For estimation, assume $w_a \sim N(\mu, \sigma^2)$. Equilibrium offers are then

$$w_i^* = \mu - \sqrt{\frac{\pi\sigma^2}{2}}, \qquad w_p^* = \mu + \sqrt{\frac{\pi\sigma^2}{2}}. \tag{7}$$

2.2 Why the Baseline Cannot Explain the Data

In our data, providers prevail in 84 percent of disputes. Equation (5) says the baseline model predicts 50 percent. We consider the two candidate explanations in turn and show that neither works, which motivates the mechanism we introduce in Section 2.3.

2.2.1 Asymmetric Risk Aversion

Prior work has shown that, in principle, differences in risk aversion can generate asymmetric win rates in final-offer arbitration (Ashenfelter and G. B. Dahl, 2012). We derive this case fully in Appendix Section A.4 and summarize three conclusions here.

First, risk aversion generates asymmetry only through the gap between the two parties' coefficients, not through its level. Under constant absolute risk aversion the equilibrium win probability is

$$P \;=\; \frac{g(\gamma_i)}{g(\gamma_i)+g(\gamma_p)}, \qquad g(\gamma) \equiv \frac{e^{\gamma\Delta}-1}{\gamma}, \quad \Delta \equiv w_p - w_i, \tag{8}$$

so that $\gamma_i = \gamma_p$ implies $P = 1/2$ however risk averse the parties are, and $\gamma_i, \gamma_p \to 0$ returns the baseline of Section 2.1. Because g is strictly increasing, the more risk-averse party is selected more often: for small coefficients, $P \approx \frac{1}{2} - (\gamma_p - \gamma_i)\Delta/8$. The relevant magnitudes are the dimensionless products $\gamma_i\Delta$ and $\gamma_p\Delta$: the level cannot create an asymmetry, but it does govern how large one is.

Second, the required asymmetry is large. Matching the observed 84 percent provider win rate requires $g(\gamma_p) = 5.25\, g(\gamma_i)$; evaluated on the most favorable possible terms, with the insurer taken to be exactly risk neutral, this gives $\gamma_p\Delta \approx 2.73$. Any positive insurer risk aversion raises this threshold on $\gamma_p\Delta$. Constant absolute risk aversion cannot then be used to argue that firm size makes this implausible, because under CARA baseline wealth cancels exactly from both first-order conditions. Under constant relative risk aversion, where the coefficient of absolute risk aversion at wealth b is ρ/b, the condition becomes $\rho_p \approx 2.73\, b_p/\Delta$, where Δ is now measured in dollars (an offer gap of roughly \$400, being 1.9 times a median 99284 QPA near \$200) — on the order of 3,400 for a solo practice and several million for a national staffing firm, against estimates in the literature of order one to ten (Cohen and Einav, 2007; Einav, Finkelstein, Ryan, Schrimpf, and Cullen, 2013).

Third, and independent of magnitudes, the mechanism requires $\gamma_p > \gamma_i$: providers must be substantially *more* risk averse than insurers. The provider organizations generating most NSA volume are private-equity-backed staffing firms filing hundreds of thousands of

near-identical disputes, among the most perfectly diversified participants in the system. The explanation requires them to be far more averse to a few hundred dollars of variance than UnitedHealthcare is. We conclude that asymmetric risk preferences cannot plausibly account for the observed win rates.

2.2.2 Arbitrator Bias

A natural alternative is that arbitrators simply favor providers, whether because they are captured, because they are compensated per case, or because they hold financial interests in provider groups. We show that this cannot be the explanation either, at least not on its own.

Proposition 1 (Neutrality of arbitrator preferences)**.** *In the baseline model, the equilibrium win rate is* $F(m^*) = 1/2$ *for* any *distribution* F *of arbitrator ideal points. Shifts in the location of* F *change equilibrium offers and awards but leave win rates unchanged.*

Proof. Equation (5) is obtained by subtracting the provider's first-order condition from the insurer's, and does not involve F except through the requirement that $F(m^*) = 1/2$ have a solution. Under a location shift $F(\cdot) \to F(\cdot - \Delta)$, equations (5) and (6) are satisfied at $m^* + \Delta$ with an unchanged offer gap, so both offers rise by Δ and the win rate is unchanged. □

Proposition 1 is worth stating plainly because it disciplines the entire exercise. If arbitrators tilt toward providers and both parties know it, both parties raise their offers, the midpoint rises with the arbitrator's ideal point, and the arbitrator remains indifferent at the margin. Awards rise; win rates do not. Any explanation of the 84 percent figure that rests on arbitrator preferences alone is therefore incomplete, and an explanation that additionally assumes insurers fail to anticipate the tilt is unattractive given who these parties are.

What Proposition 1 leaves open is the possibility that the two parties' *stakes* are asymmetric — that a dollar of offer does not mean the same thing to an insurer as it does to a provider. That is the possibility we now develop.

2.3 Publication of Offers and the Cost of Bidding

2.3.1 Institutional Basis

The NSA arbitration system is unusual among arbitration settings in that it is not private. CMS publishes a quarterly public use file reporting, for every dispute, the name of the health plan or issuer, the name of the provider or facility, the service code, the geography, the outcome, and — critically — *each party's final offer expressed as a percentage of the qualifying payment amount.* These are the fields we use to construct our analytic sample.

The consequence is that a final offer is a public statement of a firm's reservation price. When a national insurer offers three times the QPA for an emergency department visit in a given market, every provider group that negotiates with that insurer in that market can observe it. When a provider group asks four times the QPA, every payer it negotiates with can observe it. Revenue cycle management firms that specialize in NSA arbitration market this analysis explicitly.

We therefore depart from the baseline in one respect: bidding is costly, in proportion to what the bid reveals.

2.3.2 The Parties' Problems

Let M_i denote the insurer's in-network claim volume with providers other than the one in the dispute, and let M_p denote the provider's volume with payers other than the one in the dispute, both for the service in question. Let N denote the number of disputes in the dyad, and let θ_i and θ_p denote the rate at which a publicly revealed reservation price passes into subsequently negotiated rates. Define the per-dispute *publication loadings*

$$A \equiv \frac{M_i}{N}\,\theta_i \geq 0, \qquad B \equiv \frac{M_p}{N}\,\theta_p \geq 0. \tag{9}$$

Each loading is a book of business divided by a dispute count. This ratio is the source of the mechanism's quantitative bite, and it is precisely what the risk-aversion explanation lacked. There, the mechanism had to be large because the stakes were small relative to firm wealth. Here, the stakes are small relative to the book *that the bid informs*, which is exactly why a

small θ produces a large A.

We further assume that the QPA is the zero-information bid: because CMS computes and publishes the QPA from the insurer's own median contracted rate, an offer of exactly the QPA reveals nothing the market does not already know, while an offer above it reveals private willingness to pay. We write the insurer's precedent cost as $A\max\{w_i - 1, 0\}$.

Symmetrically, we assume the provider's ask is informative only up to the point where it remains credible. We take the ceiling to be billed charges $\bar{w}$: an ask above what the provider bills nobody reveals nothing about what it would accept. The provider's publication benefit is therefore $B\min\{w_p, \bar{w}\}$.[3]

The insurer and provider then solve

$$\min_{w_i \geq 0} \quad w_i F(m) + w_p\left[1 - F(m)\right] + A\max\{w_i - 1, 0\} \tag{10}$$

$$\max_{w_p} \quad w_i F(m) + w_p\left[1 - F(m)\right] + B\min\{w_p, \bar{w}\} \tag{11}$$

The first two terms of each line are exactly (3) and (4). The risk–reward tradeoff that is the heart of the Farber model is untouched; what changes is that one party now pays for talking and the other is paid for it.

2.3.3 Equilibrium

Proposition 2 (Equilibrium with published offers)**.** *At an interior equilibrium with* $1 < w_i^*$ *and* $w_p^* < \bar{w}$*,*

$$F(m^*) \;=\; \Pr(\textit{insurer wins}) \;=\; \frac{1 - A + B}{2} \tag{12}$$

$$w_p^* - w_i^* \;=\; \frac{1 + A + B}{f(m^*)} \tag{13}$$

Proof. The first-order conditions of (10) and (11) are $F(m) + A = \frac{1}{2}(w_p - w_i)f(m)$ and

[3]Some bound is necessary rather than cosmetic. With an unbounded linear benefit Bw_p, the provider's objective diverges: as $w_p \to \infty$ the term $w_p[1 - F(m)]$ vanishes for any distribution with exponential tails while Bw_p grows without bound, so no interior optimum exists. Billed charges are the natural cap because they are the largest number the provider has already asserted publicly. In our data the equilibrium ask is far below billed charges, so the cap does not bind and the characterization below is exact.

$1 - F(m) + B = \frac{1}{2}(w_p - w_i)f(m)$. Equating the left-hand sides gives (12); adding them gives (13). □

The two loadings enter the win rate with opposite signs, and the reason is worth spelling out. The insurer, facing a cost of revealing a high willingness to pay, shades its offer downward and consequently loses more often. The provider, rewarded for revealing a high reservation price, inflates its ask — and in doing so hands the arbitrator a more extreme number, which makes the insurer's offer relatively more attractive. The provider is trading arbitration wins for bargaining leverage elsewhere. The net asymmetry is governed by $A-B$; the offer gap is governed by $A + B$.

Three implications follow immediately. First, when $A = B$ the model returns exactly to (5) and (6): the benchmark requires symmetry of exposure, not the absence of externalities. Second, because $A - B$ and $A + B$ are recovered from two distinct moments — the win rate and the offer gap — both loadings are point-identified without normalization. Third, since $B \geq 0$, the win rate alone delivers the bound $A \geq 1 - 2F(m^*)$.

2.3.4 Exposure

Define each party's *NSA exposure* as the share of its business for this service that rides on arbitration, $e_i \equiv N/(N+M_i)$ and $e_p \equiv N/(N+M_p)$, so that $M/N = (1-e)/e$. Substituting into (12) with a common leak rate θ:

Proposition 3 (Exposure asymmetry)**.** *Provided an interior equilibrium exists,*

$$\Pr(\textit{insurer wins}) \;=\; \frac{1}{2} \;-\; \frac{\theta}{2}\left(\frac{1}{e_i} - \frac{1}{e_p}\right). \tag{14}$$

When the interior equilibrium exists

Equation (14) is derived from both parties' first-order conditions and therefore presupposes that both have interior solutions. Since a probability must lie in $[0, 1]$, this requires

$$|A - B| \;\leq\; 1, \qquad \text{equivalently} \qquad \theta\left|\frac{1}{e_i} - \frac{1}{e_p}\right| \;\leq\; 1. \tag{15}$$

The restriction is not a technical nuisance; it marks the boundary between the two regimes the model describes. When $A - B > 1$ the insurer's marginal cost of raising its offer exceeds the marginal benefit at every offer level, its first-order condition has no root, and it sits at the corner of Section 2.4 — that is, it submits a protest bid. The observed win rate in that case is governed by the mixture expression (18), not by (14). Exposure asymmetry beyond the threshold in (15) therefore does not push the interior win rate below zero; it moves the insurer out of the interior game altogether.

Two implications are worth noting. First, the interior mechanism can rationalize any insurer win rate in $[0, 1]$, with rates above one half when the provider is the more exposed party ($B > A$), as in Appendix Figure 1; it is falsifiable because an insurer that contests seriously and still never wins cannot be accommodated. Second, our estimate $\widehat{A - B} = 0.669$ sits comfortably inside (15), so the mechanism is not being stretched to its limit to fit the data. For illustration, an insurer with exposure $e_i = 0.1$ facing a fully exposed provider ($e_p = 1$) matches the observed 16.6 percent win rate (Table 3, interior disputes, $N = 303{,}533$)[4] at $\theta \approx 0.074$: a seven percent pass-through, magnified ninefold by the exposure gap.

Equation (14) is the central result. The insurer wins less than half the time exactly when it is less NSA-exposed than the provider. The emergency medicine staffing groups that dominate NSA arbitration concentrate their revenue in precisely the codes that go to arbitration, so e_p is large; any national insurer has an in-network book that dwarfs its dispute volume, so e_i is near zero. The gap is enormous, and the provider wins.

The same equation predicts where this should fail. When both parties are diversified — a large hospital system disputing with a regional insurer — the bracket vanishes and the baseline model holds exactly. Appendix Figure 1 documents this for the largest such dyad.

2.4 Protest Bids as a Corner Solution

Because the insurer's precedent cost applies only to offers above the QPA, its objective (10) is kinked at $w_i = 1$. This generates a mass of insurers bidding exactly the benchmark rather

[4]Table A1 reports 15.8 percent on the 560,476-dispute large-provider subsample used in Table 1; the full emergency-department sample is 596,608 disputes. Restricting to CPT 99284 gives $\widehat{A - B} = 0.659$.

than a smooth distribution around it.

Proposition 4 (Protest condition)**.** *Let* $m_1 \equiv (1 + w_p)/2$*. The insurer's best response is* $w_i^* = 1$ *exactly if*

$$A \;\geq\; \tfrac{1}{2}(w_p - 1)f(m_1) \;-\; F(m_1). \tag{16}$$

Proof. $w_i = 1$ is optimal if and only if the left derivative of the objective at 1 is non-positive and the right derivative is non-negative. The left derivative omits A and the right derivative includes it, which yields (16). □

We refer to offers at this corner as *protest bids*. The interpretation is not that the insurer is behaving irrationally or failing to anticipate the arbitrator. It is that contesting the dispute would require publicly revealing a willingness to pay above the statutory benchmark, and for an insurer with a large in-network book that revelation is more expensive than the dispute is worth. The insurer offers the QPA, loses, and pays. In the data, 36.4 percent of insurer offers are exactly the QPA and a further 1.9 percent are zero; the share ranges from 0.2 percent for BlueCross BlueShield of Tennessee to 90 percent for Aetna (Table 3). A smaller residual group offers zero, which we interpret as insurers additionally contesting the eligibility of the claim.

Protest bids in turn change the provider's problem. Let q denote the probability the insurer is at the corner. Since the provider submits its offer without observing the regime, it solves

$$\max_{w_p} \; q\left[F(m_1) + w_p\left(1 - F(m_1)\right)\right] + (1 - q)\left[w_i^* F(m) + w_p\left(1 - F(m)\right)\right] + Bw_p. \tag{17}$$

Proposition 5 (Bid inflation against protesters)**.** $dw_p^*/dq > 0$ *whenever* $f(m_1) < f(m^*)$*.*

The intuition is that in the protest regime the insurer's offer is pinned at the QPA and does not respond to the provider's ask, so the only cost of asking more is the thin probability that the arbitrator's ideal point falls below $(1 + w_p)/2$. Facing an insurer that protests, the provider asks for a great deal and accepts losing occasionally. Facing an insurer that contests, the provider moderates. This is a sharp prediction about the joint distribution

of offers within insurer–provider pairs, and Table 3 tests it.

Combining the two regimes, the unconditional provider win rate is

$$\Pr(\text{provider wins}) \;=\; q\,[1 - F(m_1)] \;+\; (1-q)\,\frac{1 + A - B}{2}. \tag{18}$$

The first term exceeds one half, and the second does whenever $A > B$. This is why removing protest bids from the sample does not restore balanced win rates: the asymmetry operates on the intensive margin as well as at the corner.

2.5 Arbitrator Inducement and Regulatory Capture

Proposition 2 determines who wins, but it takes the location of F as given. By Proposition 1, that location is exactly what does *not* affect win rates — but it is what determines award levels, and award levels are what ultimately matter for health care prices. This is the division of labor in our model: exposure asymmetry explains who wins, and arbitrator incentives explain how much they win.

Under the NSA, arbitration entities are compensated on a per-case basis. This creates an incentive for arbitrators to consider not only the outcome of the current dispute but the effect of their decisions on future case volume. We model the arbitrator as maximizing

$$U_a = -\alpha(x - w)^2 + (1 - \alpha) f_a n, \qquad n = k(w - x), \tag{19}$$

where x is the arbitrator's ideal outcome absent inducement, f_a is the fee per case, k parameterizes how strongly award levels induce case flow, and $\alpha \in [0, 1]$ is the weight on deciding the case on its merits. Solving yields a decision standard

$$\mu_a \;=\; x \;+\; \frac{(1-\alpha) f_a k}{2\alpha} \;+\; \gamma \omega_a, \tag{20}$$

where we have added a third term allowing the arbitrator's financial interest in provider-side surplus, ω_a, to shift the standard directly.

Equation (20) makes explicit a point that levels alone cannot settle: inducement and

direct capture both raise μ_a, and a cross-sectional relationship between arbitrator identity and award levels does not distinguish them. What does distinguish them is dynamic. Inducement implies that an arbitrator's decision standard should predict its *subsequent case volume*, and that the standard should rise with the fee. Direct capture implies a level shift with no volume response once ownership is controlled for. We implement this test in Section 4.

A final observation connects this section to the previous one. Both disputing parties participate in selecting the arbitration entity and both may object to a selection, which would seem to preclude systematic tilt toward either side. But on the substantial share of disputes where the insurer is at the corner of Proposition 4, the insurer is not contesting the dispute and exercises no meaningful selection pressure. On those cases selection is effectively unilateral, and arbitration entities competing for case flow face a provider-tilted marginal incentive whether or not any ownership link exists. Documented financial ties between arbitration entities and provider groups are the sharpest instance of a force that competition for selection generates generically.

2.6 From Arbitration to Negotiated Prices

The preceding sections describe behavior within a dispute. But arbitration under the NSA is not only a dispute-resolution mechanism; it is an outside option that governs what happens when an insurer and a provider fail to agree on an in-network contract. This section makes that link explicit and generates the prediction we take to price data.

Consider an insurer and a provider bargaining over an in-network rate r. Let $s \in [0, 1]$ denote the share of the provider's volume with this insurer that is NSA-eligible: emergency and ancillary services where the patient does not choose the provider. On the remaining $1 - s$ of volume — scheduled and elective care — remaining out of network yields the provider approximately nothing, because patients substitute toward in-network providers. Disagreement values per unit of volume are therefore

$$v_p = s\left(\mathbb{E}[w] - \kappa_p\right), \qquad d_i = s\left(\mathbb{E}[w] + \kappa_i\right) + (1 - s)c_0, \tag{21}$$

where κ_p and κ_i are per-case arbitration costs and c_0 is the insurer's cost of routing patients to an alternative in-network provider. Nash bargaining with provider weight β gives $r^* = \beta d_i + (1-\beta)v_p$:

Proposition 6 (Pass-through of awards into negotiated rates)**.**

$$r^* = s\,\mathbb{E}[w] + \beta s\kappa_i - (1-\beta)s\kappa_p + \beta(1-s)c_0, \qquad so \qquad \frac{\partial r^*}{\partial \mathbb{E}[w]} = s. \tag{22}$$

Three testable implications follow. First, post-NSA growth in negotiated in-network rates should be increasing in the provider's NSA-exposed share s: providers who use arbitration should see their negotiated prices rise relative to those who do not. Second, the effect should concentrate in NSA-heavy service codes and be absent for scheduled care. Third, the model predicts a specific and otherwise puzzling time path around implementation. Balance billing was prohibited in January 2022, but the arbitration process was stayed by litigation for much of that year. During that window $\mathbb{E}[w]$ was near zero and (22) implies negotiated rates should have *fallen* for exposed providers, before rising above baseline once arbitration came online and $\mathbb{E}[w]$ exceeded the pre-NSA value of balance billing. We show in Section 5 that this is what the data do.

Finally, note that s in (22) and e_p in (14) are the same underlying object: how concentrated a provider is in services the NSA governs. The parameter that determines who wins the arbitration is therefore the parameter that determines how much negotiated prices move, and the two estimates must agree. This is an overidentifying restriction we intend to test as the exposure data are built.

3 Data

3.1 NSA Arbitration Data

We use publicly available CMS arbitration data from over a million disputes between the start of the law in 2023 and the end of 2024. The data elements include (a) final offers from both parties (i.e., providers and insurers), (b) which party prevailed, (c) payer name,

(d) provider name, (e) CPT, and (d) location (MSA). We make a few restrictions, such as omitting default judgments (where one party wins because another does not submit required paperwork) and omitting early quarters of the data where court cases pause and started arbitration, rendering our confidence in that data lower.[5] In some analyses, we also restrict to the largest insurers and providers by number of disputes and to the top 3 CPTs to focus on the most frequent users of arbitration. Although we focus on the top CPT codes, which are all for emergency room care, our results are robust to looking at all CPTs.

In the CMS public files, the dollar amounts of each dispute are held in different files than the payer and provider information and—with limited exception— they cannot be merged. However, we can still back out levels of payment without merging due to an element of the law called the Qualifying Payment Amount (QPA). The No Surprises Act legally requires arbitrators to consider several factors when choosing between offers. These include provider quality, experience, and teaching status, market share of each party, a prohibition on considering Medicare rates and considering provider charges, and the qualifying payment amount (QPA). This amount is an insurer-specific median of contracted rates for a service, insurance type, and geography in 2019, adjusted for inflation. We report the offer amounts relative to the QPA in our analysis.

The top providers in arbitration are TeamHealth, SCP Health, and Envision. Many of the frequent users of the IDR process are private equity-backed emergency room staffing firms who had large shares of business in out-of-network billing prior and wielded significant price leverage prior to the No Surprises Act (Biener, Chartock, Garmon, and Trish, 2021; Cooper, Scott Morton, and Shekita, 2020).

3.2 Other Sources of Data

We augment our primary data source with actual negotiated price data released under the Transparency in Coverage (TiC) rule starting in 2022. This rule mandated that all payers post their negotiated price for each Provider–Procedure–Plan. We use these prices

[5] A lawsuit filed by the Texas medical association (*TMA*) was the main source of this pause and slowdown. https://www.theguardian.com/us-news/2023/sep/29/no-surprises-act-balance-billing-doctors-insurance-companies

to compare arbitration outcomes to insurer-provider negotiated in-network rates nationally. We focus specifically on prices for standard emergency medicine professional services (CPT code 99284). The TiC data have previously been used to compare geographic and insurer variation for these services (Philips and C. M. Whaley, 2025).

3.3 Summary Statistics

Arbitrators are awarding providers wages that are higher than what insurers are paying for the same care under in-network contracts. Figure 1 plots the 25th, 50th, 75th, and 90th percentile of arbitration awards from the CMS data relative to the distribution of negotiated rates from Transparency in Coverage data, focusing on one representative service, an emergency room visit of mid-level severity (CPT 99284), the most common service code disputed under the No Surprises Act. Nearly all of the negotiated rates are lower than the 25th percentile of the arbitration decisions (only 4% are higher), and only 1.5% of negotiated rates exceed the median award. Superimposed on this graph is also the QPA amount from the CMS data, which is slightly below the median negotiated rate. In this figure the mean NSA awarded rate is 3.2 times higher than mean in-network rates ($736 vs. $231).

Appendix Table 1, Panel A shows that providers win frequently in arbitration. In our main sample, 84% of arbitration disputes are won by providers. In Appendix Table 1 Panel B, we show offers from providers are at the median 2.97 times the QPA.[6] These results are consistent with other descriptive papers (e.g., (Adler and Fiedler, 2024)).

Figure 2 shows the relationship between offers of the insurer (x-axis) and offers of the provider (y-axis). This scatterplot is shown for the most frequent pairing of disputants, the UnitedHealthcare vs. TeamHealth dyad. By construction, all of the data is above the 45° line. If the final offer of the provider was lower than the final offer of the insurer, they should agree to settle without spending resources on arbitrator fees.[7] A primary objective of this figure is to highlight that offers from a given party (insurer or provider) are equilibrium objects: a higher offer from one party impacts the other side's offer. Displaying these

[6] Slight differences between this and Figure 1 relate to Figure 1 coming from CMS's dollar file dataset and Table 1 from the relative-to-QPA dataset.

[7] In making this figure, we drop a small number of off-equilibrium offers where this inequality does not hold.

together shows this relationship and the tradeoff on the margin between additional risk of losing and more income earned from an aggressive final offer.

This figure also demonstrates that insurers bid exactly the QPA in 36 percent of disputes and zero in 2 percent. A zero bid could be rationalized by the baseline model alone—if the insurer thought the arbitrator's ideal point was 2 times the QPA and the provider bid more than 4 times, bidding zero is a best response—but the mass at exactly the QPA cannot. Section 2.4 shows that an offer at the QPA is the corner of the insurer's problem when its in-network exposure is large, and Table 3 reports its incidence by insurer.[8]

4 Results

4.1 Protest bids and the provider's response

Table 3 reports where insurers put their offers. Across 596,608 emergency-department disputes, 36.4 percent of insurer offers are exactly the QPA and a further 1.9 percent are exactly zero; only 6.7 percent lie strictly between. The mass at the benchmark is a firm-level choice. BlueCross BlueShield of Tennessee offers the QPA in 0.2 percent of its disputes and UMR in 4 percent, while UnitedHealthcare does so in 29 percent, Anthem in 77 percent, and Aetna in 90 percent. UMR is the clearest case of an insurer that contests at the minimum: 93 percent of its offers lie in $(1, 1.5]$. Providers win 85.6 percent of protested disputes and 83.5 percent of contested ones.

Proposition 5 predicts that a provider asks more of an insurer that protests more. The bottom panel of Table 3 regresses the log provider ask on the share of the counterparty insurer's other offers in the same CPT code and quarter that are protests. The coefficient is 0.25 (standard error 0.10, clustered by insurer) with provider-group fixed effects and 0.32 (0.15) without: a provider facing an insurer that always protests asks roughly 25 percent more than one facing an insurer that always contests. The comparison is across insurers; within an insurer, contested disputes are selected on high stakes, so we do not use within-

[8] Insurers also contest the eligibility of many claims outright; lawsuits over non-payment proliferate (`https://litigationtracker.law.georgetown.edu/issues/no-surprises-act/`).

insurer variation.

On disputes in which the insurer offers above the QPA and below the provider ($N = 303{,}533$), the insurer wins 16.6 percent, which by equation (12) gives $\widehat{A - B} = 0.669$. The unconditional rate over all 596,608 disputes is also 16.6 percent, and restricting to CPT 99284 gives 0.659.

4.2 Arbitrator decision standards

In this section we show results from our baseline structural model and structural parameters from the inducement model. Our regressions show high estimates of the center of the arbitrator's preferred distribution of arbitration award relative to QPA as well as a wide spread around that award. Our primary specification restricts to the third quarter of 2023, and finds a mean of 10.7 and a standard deviation of 9.7 times the QPA (Table 1, column 1). [9]

We interpret these estimates with skeptical caution. The Farber model of final offer arbitration, upon which we base our structural estimates, suggests that under symmetric risk aversion, it is equally as likely for either side to win (see Mas, 2006). Given the lopsided win rate in our dataset, we acknowledge that parameters of μ and σ are estimated based on support in the data mostly from the left tail of the estimated distribution. Section 4.1 accounts for the win rate. Table 2 turns to the *level* of awards, which by Proposition 1 is where arbitrator incentives operate.

In Table 2, we calibrate the inducement model to fit the data used in the earlier probit regressions by estimating an arbitrator-specific α, backing out the degree of inducement using the variation in arbitrators and fees. We follow similar methodolgy matching historical fees developed by Valtadoros, 2026. We find arbitrators vary in their degrees of inducement, but all demonstrate a tendency to diverge meaningfully from the QPA when the incentive for induced revenue is incorporated in their utility function. The details of this estimation are found in the Appendix. In the Appendix, we also document that providers with a

[9] Other specifications show very similar estimates when excluding fixed effects for different emergency CPT codes and an indicator for whether the case was batched, or rolled up surprise bill disputes for the same code X provider X insurer.

greater concentration of services exposed to the NSA are more likely to initiate cases. Other providers, such as large hospital systems with relationships with insurers that cover a broad range of non-NSA services, are less likely to initiate cases. [10]

In Table 2, we find meaningful variation in our α estimates, signaling variation in IDRE inducement.[11] For example, the "Island Peer Review" IDRE has the highest μ' and among the lowest α estimates, while "Keystone Peer Review" and "Provider Resources" have the lowest estimated α values, consistent with inducement of cases by IDREs. We also note that many IDREs with low estimated α parameters have either direct or explicit ownership or financing connections to private equity companies, many of which also have investment stakes in provider groups bringing NSA cases.[12] These apparent conflicts raise questions about the independence of IDREs. We also note that both the initial selection and the "veto" ability of IDRE participants introduces competition between IDREs. Even if an IDRE is not sponsored by a firm that may also have an ownership stake in providers bringing IDRE cases, competition with potentially conflicted IDREs may lead to similar inducement and decision responses.

5 Impact of NSA Arbitration on In-Network Prices

Our previously-estimated model using NSA arbitration data shows providers prevail in approximately 84% of cases and are awarded rates of approximately 3× QPA. Consistent with insurer-provider bargaining models, the "outside option" of NSA arbitration may have an impact on provider-insurer network formation and negotiations. Thus, despite the NSA's policy focus on explicitly non-negotiated and out-of-network payment, there could be an

[10] We also compare provider win rates in New York State and provider win rates in the NSA separately for two out of the three arbitrators who conducted arbitration disputes both in New York State prior to the passage of the NSA as well as under the NSA using Freedom-of-Information records from New York State. (The third arbitrator who worked in New York State does not appear to be involved in NSA disputes.) In New York prior to the NSA, there was a law requiring the arbitrators to consider the eightieth percentile of billed charges—a number higher than the QPA. When these same two arbitrators resolve disputes under the NSA, they select the provider's offer far more frequently than the other arbitrators, perhaps because they have experience with high provider reimbursement already. These two firms choose the provider 96% of the time (in over 43,000 disputes), which is well higher than the baseline provider win rate of 84%.

[11] Our model is identified ordinally, not cardinally, as the parameter k is a floating parameter which we set equal to 0.8 but shifts around α.

[12] Source: https://pestakeholder.org/news/profiting-on-all-sides-private-equity-and-the-no-surprises-act/

indirect impact on in-network participation and prices.

To estimate this impact, we use data from a large national insurer. Among this insurer, we focus on all claims for the most common NSA service, emergency medicine physician care (CPT code 99284). We extract all 99284 claims over the 2020-2025 period, resulting in 1.6 million claims across approximately 50,000 providers. Approximately 2% of claims are among out-of-network providers. Among these providers, we construct a quarterly panel on the provider's mean negotiated reimbursement rate (among in-network providers) and the share of out-of-network claims.

We link these data to the data on NSA arbitration events. For each provider, identified and linked between the two datasets by their unique National Provider Identifier (NPI), we identify if the provider ever engages in NSA arbitration with the national insurer. Among those with NSA cases, we also identify the date of their first arbitration. With this panel and linkage, we estimate an event study of the form:

$$y_{it} = \alpha + \sum_{\substack{j=-24 \\ j\neq -1}}^{12} \beta_j \mathbf{1}[NSA_j \times Time_t = j] + \lambda_i + \tau_t + \varepsilon_{it} \tag{23}$$

where Y_{it} is the outcome variable for provider i in quarter t, and $1[NSA_i \times Time = j]$ represents a series of indicators equal to 1 if the provider has engaged in NSA arbitration and relative time is equal to j, λ_i are provider fixed effects, and τ_t are quarter fixed effects. β_j are the coefficients of interest. We cluster standard errors at the provider level. To account for potential bias due to differential treatment timing and heterogeneous treatment effects, we apply the Callaway and Sant'Anna estimator (Callaway and Sant'Anna, 2021).

5.1 In-Network Results

Figure 4 presents descriptive trends that underlie our empirical approach. Within this figure, the orange line presents in-network price trends among providers that never engaged with the NSA arbitration process, while the blue line presents trends among those with any NSA arbitration case. Price trends among providers not exposed to the NSA are stable and increase steadily over the 2020 to 2025 period. In contrast, price trends for providers that

ever engage with the NSA process are higher in the pre-NSA period. Reimbursement rates fall in the immediate NSA implementation period (2022), but then steadily increase and diverge from the non-NSA provider price trends.

Figure 5 presents the event study version of this dynamic. In the periods prior to a provider's first NSA case, in-network price trends are stable. Following a provider's first NSA case, there is an immediate increase in in-network prices, which stabilizes around a 20% increase. The average treatment effect is 0.23 log-points, equivalent to a 27% increase in negotiated rates. This result is consistent with findings by Prager and Tilipman, 2020 and illustrate how increased outside-options, in this case induced by the NSA's arbitration process, can lead change negotiation leverage in insurer-provider negotiations. In this case, this increased negotiation leverage is used to negotiate substantially higher in-network rates.

6 Conclusion

In this paper, we find that the arbitration system of the No Surprises Act leads to high awards for providers and high provider win rates, and we explain this with a micro-founded model of arbitrator behavior. We introduce economic incentives by both parties that may lead to the observed outcomes and estimate models of arbitrator preferences. A standard model of arbitration is not consistent with the observed skewness in outcomes. Exposure asymmetry from published offers accounts for who wins; the incentives of arbitrators themselves account for how much they win.

These arbitrator incentives, which in our model set the level of awards, are critical to consider when forecasting the impact of the NSA on US health care markets and prices. Prior to the NSA, providers and insurers negotiated prices that are based on network formation and reflect consumer demand (Katherine Ho, 2009; Gowrisankaran, Nevo, and Town, 2015; Kate Ho and Lee, 2019). The current implementation of the NSA creates a feasible "outside option" for providers to obtain higher payment rates without the countervailing pressure of insurer's restricting access to patients.

Based on our data, providers with exposure to the NSA may make the calculation of

going through the NSA process with a win rate of approximately 84 percent and a payoff of approximately 3 times the QPA when it wins and roughly the QPA when it loses, and thus an expected payoff of 2.7 times the QPA, a value which we show is accurately close to median in-network rates. For the most common NSA service, standard emergency medicine professional services (CPT code 99284), mean in-network negotiated rates are approximately $250, compared to the median rate of $200. Thus, for this single service, going through the NSA process results in an expected payoff of $540, or a 115 percent increase over prevailing in-network rates. Nationally, approximately 9 million in-network 99284 services occur per year, at a total cost of $2.3 billion. At the upper extreme and assuming no changes to the NSA process, providers in aggregate may capture an additional $2.7 billion for just one service, either in the form of NSA award amounts or increased payments to remain in-network. Extrapolating beyond this single procedure to all emergency medicine procedures results in substantially higher estimates.

As earlier hypothesized (Ashenfelter, Bloom, and G. Dahl, 2013), we see asymmetric use of expert agents in arbitration with providers using consultants more frequently.[13] Secondly, we would expect that firms with more at stake because most of their business is exposed to NSA services (the providers) would submit more paperwork and documentation than insurers to arbitrators. This has been confirmed in conversations with CMS administrators and arbitration participants–providers are "flooding the zone" with paperwork that insurers are simply not submitting, and they are far more frequently the initiating party of the dispute.

Several policies can ameliorate concerns that the NSA is leading to health care price inflation. First, arbitrators can consider publicly available Transparency-in-Coverage data when arriving at dispute conclusions. Price transparency data provides the distribution of negotiated prices, allowing an arbitrator to sidestep any potential controversy surrounding the QPA and its method of calculation. It is also similar to what Washington State uses for their state-level arbitration system. Washington had less than 150 arbitration cases per

[13]Nutex Health, one of the revenue cycle management consultants, reported in a recent 10-K that their total revenue was "driven by their successful participation in arbitration through the Independent Dispute Resolution Process of the No Surprises Act". https://earningscall.biz/e/nasdaq/s/nutx/y/2024/q/q4

year prior to the NSAs passage. This suggests that use of price transparency data can lower the incidence of costly arbitration disputes.

Second, regulators can hold a scoring auction for arbitrator procurement. In light of our model and suggestive evidence of the incentives for arbitrators to induce more cases, CMS could design a scoring auction that explicitly certifies arbitrators for a limited time with re-accreditation based on factors such as use of price distribution data and speed of resolving cases. If the arbitrators were unable to meet certain policy objectives, they could run the risk of scoring low and not being selected to hear disputes. It seems arbitrator competition has been less than fully effective.

And finally, considering restrictions on vertical integration between arbitrators and the parties using arbitration may lead to less inducement and is consistent with existing rules of the NSA surrounding arbitrator eligibility. Recent reports suggest vertical integration between providers using arbitration and the arbitrators themselves is more widespread than previously known.

Arbitration incorporates the theoretical benefit of forcing one party to internalize the other party's bidding decision when forming their own bid. Harnessing market dynamics when there was no negotiated agreement to begin with captures the spirit of incomplete contracts *ex-ante* while limiting holdup *ex-post*. However, arbitration as currently implemented results in high payouts to providers and potentially greater payments than they earned prior to the No Surprises Act when they could simply balance bill patients directly. Frequent public dissatisfaction with arbitration is a symptom of a mechanism whose published offers and per-case arbitrator pay were not designed with these incentives in mind.

Finally, we close by noting that one can also view arbitration as an insurance system itself: patients are fully insured from out-of-network surprise bills, but at a loading cost of the increase in price caused by upward pressure on prices. The arbitration procedure induced a welfare tradeoff in that after it is introduced, patients can no longer get surprise bills, yet the negotiated prices that every enrollee faces may be higher due to increased leverage in negotiations coming from a high outside option of using arbitration relative to negotiating an in-network rate. More research is needed to further understand this welfare

tradeoff and we anticipate learning more as data on network membership and prices becomes available.

References

Adler, Loren and Matthew Fiedler (2024). "Outcomes under the No Surprises Act arbitration process: A brief update". In.

Ashenfelter, Orley and David Bloom (1984). "Models of Arbitrator Behavior: Theory and Evidence". In: *The American Economic Review* 74.1, pp. 111–124. ISSN: 00028282. URL: http://www.jstor.org/stable/1803312.

Ashenfelter, Orley, David Bloom, and Gordon Dahl (2013). *Lawyers as agents of the devil in a prisoner's dilemma game: Evidence from long run play*. Tech. rep. National Bureau of Economic Research.

Ashenfelter, Orley and Gordon B Dahl (2012). "Bargaining and the role of expert agents: An empirical study of final-offer arbitration". In: *Review of Economics and Statistics* 94.1, pp. 116–132.

Biener, Adam I, Benjamin Chartock, Christopher Garmon, and Erin Trish (2021). "Emergency Physicians Recover A Higher Share Of Charges From Out-Of-Network Care Than From In-Network Care: Study examines how much privately insured emergency department patients paid when they likely received a surprise bill and how much physicians received in these situations." In: *Health Affairs* 40.4, pp. 622–628.

Callaway, Brantly and Pedro H.C. Sant'Anna (2021). "Difference-in-Differences with multiple time periods". In: *Journal of Econometrics* 225.2. Themed Issue: Treatment Effect 1, pp. 200–230. ISSN: 0304-4076. DOI: https://doi.org/10.1016/j.jeconom.2020.12.001. URL: https://www.sciencedirect.com/science/article/pii/S0304407620303948.

Cohen, Alma and Liran Einav (2007). "Estimating risk preferences from deductible choice". In: *American economic review* 97.3, pp. 745–788.

Cooper, Zack, Fiona Scott Morton, and Nathan Shekita (Sept. 2020). "Surprise! Out-of-Network Billing for Emergency Care in the United States". In: *Journal of Political Economy* 128.9. Publisher: The University of Chicago Press, pp. 3626–3677. ISSN: 0022-3808.

DOI: 10.1086/708819. URL: https://www.journals.uchicago.edu/doi/full/10.1086/708819.

Einav, Liran, Amy Finkelstein, Stephen P Ryan, Paul Schrimpf, and Mark R Cullen (2013). "Selection on moral hazard in health insurance". In: *American Economic Review* 103.1, pp. 178–219.

Farber, Henry S (1980). "An analysis of final-offer arbitration". In: *Journal of conflict resolution* 24.4, pp. 683–705.

Garmon, Christopher and Benjamin Chartock (2017). "One in five inpatient emergency department cases may lead to surprise bills". In: *Health Affairs* 36.1, pp. 177–181.

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town (2015). "Mergers when prices are negotiated: Evidence from the hospital industry". In: *American Economic Review* 105.1, pp. 172–203.

Grossman, Sanford J and Oliver D Hart (1986). "The costs and benefits of ownership: A theory of vertical and lateral integration". In: *Journal of political economy* 94.4, pp. 691–719.

Gupta, Atul, Sabrina T Howell, Constantine Yannelis, and Abhinav Gupta (2024). "Owner incentives and performance in healthcare: Private equity investment in nursing homes". In: *The Review of Financial Studies* 37.4, pp. 1029–1077.

Gupta, Atul, Ambar La Forgia, and Adam Sacarny (2024). *Turbocharging Profits? Contract Gaming and Revenue Allocation in Healthcare*. Tech. rep. National Bureau of Economic Research.

Ho, Kate and Robin S Lee (2019). "Equilibrium provider networks: Bargaining and exclusion in health care markets". In: *American Economic Review* 109.2, pp. 473–522.

Ho, Katherine (Mar. 2009). "Insurer-Provider Networks in the Medical Care Market". In: *American Economic Review* 99.1, pp. 393–430. DOI: 10.1257/aer.99.1.393. URL: https://www.aeaweb.org/articles?id=10.1257/aer.99.1.393.

Hoadley, Jack and Kennah Watts (Aug. 2025). "The Substantial Costs Of The No Surprises Act Arbitration Process". en. In: *Health Affairs Forefront*. DOI: 10.1377/forefront.

20250820.13433. URL: https://www.healthaffairs.org/do/10.1377/forefront.20250820.13433/full/ (visited on 10/26/2025).

Kong, Yunmi, Bernardo S Silveira, and Xun Tang (2025). "Risk and information in dispute resolution: An empirical study of arbitration". In: *Journal of Political Economy* 133.9, pp. 2794–2835.

Mas, Alexandre (2006). "Pay, reference points, and police performance". In: *The Quarterly Journal of Economics* 121.3, pp. 783–821.

Philips, Alexander P. and Christopher M. Whaley (2025). "Commercial Price Variation for Emergency Services Evaluation and Management". In: *Annals of Emergency Medicine* 86.3, pp. 296–298. ISSN: 0196-0644. DOI: https://doi.org/10.1016/j.annemergmed.2025.04.016. URL: https://www.sciencedirect.com/science/article/pii/S0196064425002100.

Prager, Elena and Nicholas Tilipman (2020). *Regulating Out-of-Network Hospital Payments: Disagreement Payoffs, Negotiated Prices, and Access*. Tech. rep. Working Paper.

— (Apr. 2025). *Regulating Out-of-Network Hospital Payments: Disagreement Payoffs, Negotiated Prices, and Access*. Working Paper 33727. National Bureau of Economic Research. DOI: 10.3386/w33727. URL: http://www.nber.org/papers/w33727.

Singh, Yashaswini (Oct. 2022). *Can Private Equity Buy Referrals? Evidence from Multispecialty Physician Practice Acquisitions*. en-us. Working Paper. URL: https://www.yashaswinisingh.com/publication/job-market_paper/.

Singh, Yashaswini, Nandita Radhakrishnan, Loren Adler, and Christopher Whaley (2025). "Growth of private equity and hospital consolidation in primary care and price implications". In: *JAMA Health Forum*. Vol. 6. 1. American Medical Association, e244935–e244935.

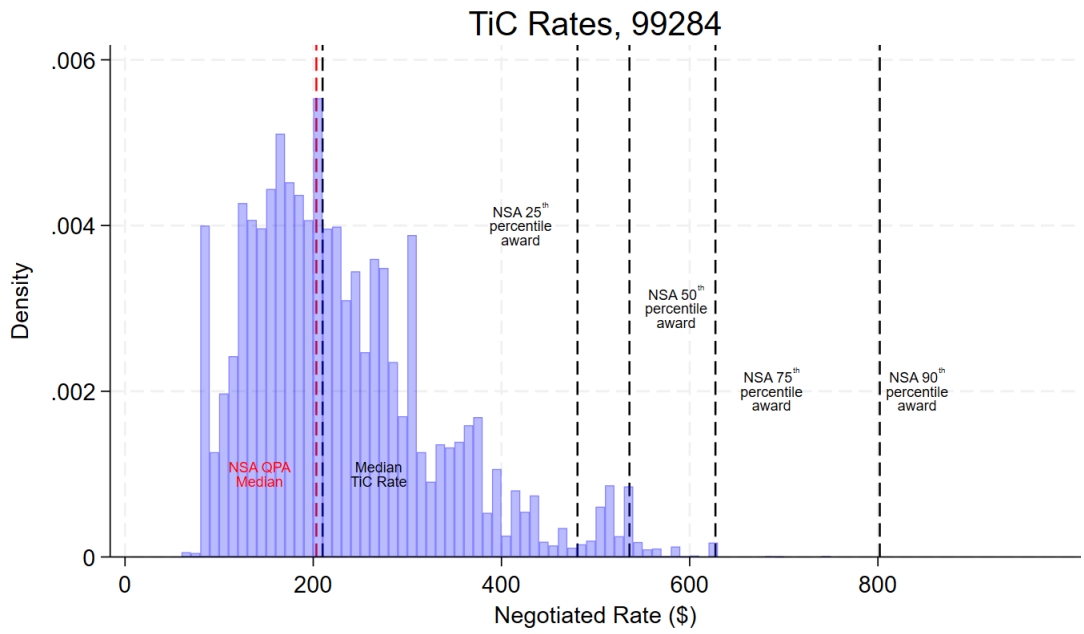
Valtadoros, Philip (2026). "Arbitrator Differentiation and Competition under the No Surprises Act". In: *Available at SSRN 6235018*.

Williamson, Oliver E (1975). "Markets and hierarchies: analysis and antitrust implications: a study in the economics of internal organization". In: *University of Illinois at Urbana-*

Champaign's Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship.

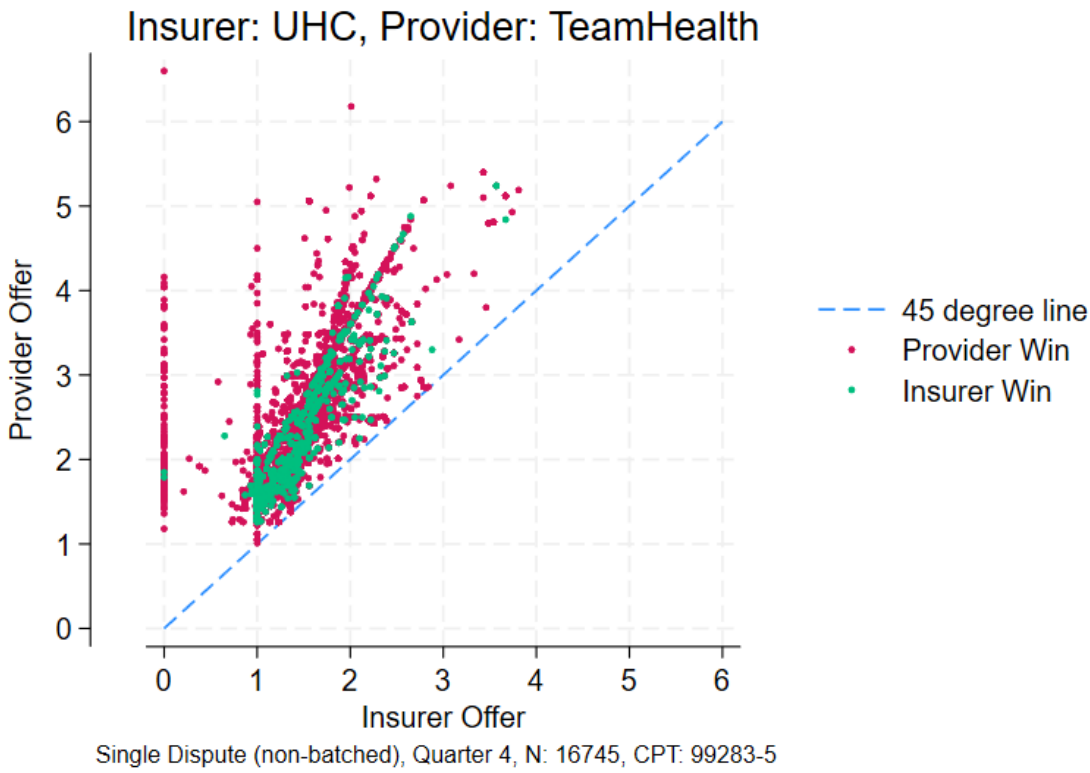
7 Tables and Figures

Figure 1: Arbitration Awards Relative to the QPA for Most Common Service (Emergency Care, Moderately Complex)



Note: Transparency in Coverage data first used in Philips & Whaley (2025). Arbitration awards relative to QPA come from CMS data. The median QPA in the NSA data is approximately equal to the median payment found in the TiC data. Only 1.5% of all negotiated rates are higher than the median NSA award.

Figure 2: Offers Relative to Qualifying Payment Amount (QPA)
Restricted to Top Insurer–Provider Dyad in the Data Set



Note: Data is mechanically above the 45-degree line as it is off-equilibrium path to have an offer from the insurer higher than the offer from the provider (as they would simply find an agreed upon rate between those two offers and save by not paying an arbitrator fee). Only a small number of cases are off-equilibrium path, which we consider to be communication breakdowns and are dropped from the analytical sample. The mass at 1.0 is the protest bid of Table 3.

Table 1: Maximum Likelihood Estimator (Probit on Insurer Victory)

	(1) inswin	(2) inswin	(3) inswin
inswin			
wage avg.	0.103***	0.112***	0.111***
	(0.0216)	(0.0195)	(0.0195)
constant	-1.099***	-1.223***	-1.357***
	(0.199)	(0.189)	(0.187)
Insurer & Prov FEs	Yes	Yes	Yes
Location FEs	Yes	Yes	Yes
CPT FEs	Yes	Yes	No
Batched?	Yes	No	No
μ	10.7	10.9	12.2
σ	9.7	8.9	9.0
Pseudo R-squared	0.344	0.340	0.338
Observations	23,536	24,891	24,891

Standard errors in parentheses

Note: Q3 Data only.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Presented here are estimates from the structural model described in Section II. Quarters 1 & 2 were not retained due to starts & stops in the law subject to litigation brought by the Texas Medical Association in federal court.

Table 2: IDRE Parameters

IDRE	μ'	α
Island Peer Review Organization	22.202	0.065
EdiPhy Advisors	21.287	0.073
C2C Innovative Solutions	17.420	0.069
iMPROve Health	15.040	0.081
Federal Hearings and Appeals Services	12.909	0.076
Maximus Federal Services	7.356	0.156
National Medical Reviews	7.291	0.117
ProPeer Resources	6.340	0.199
Medical Evaluators of Texas	5.996	0.132
Network Medical Review Company	5.922	0.140
MCMC Services	4.116	0.298
Keystone Peer Review Organization	-1.536	-0.898
Provider Resources	-1.696	-0.590

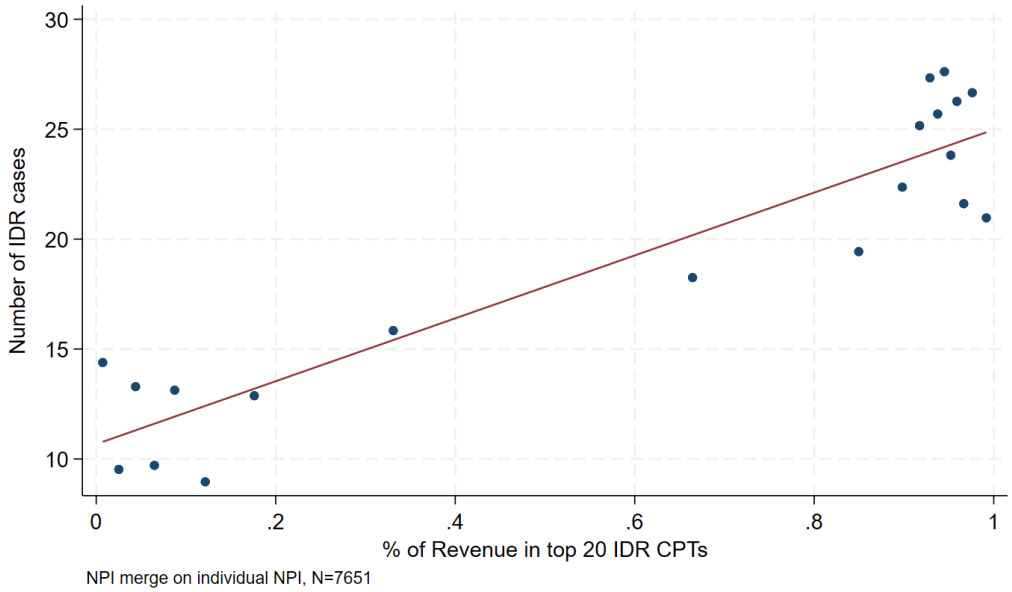
Note: Data on IDRE is merged to CMS data through arbitrator fee and year matching as described in Appendix Section 8.1. Due to free-floating parameter k in the model, α is identified ordinally, not cardinally.

Table 3: Protest bids by insurer, and the provider's response

		Share of insurer offers, relative to the QPA				
Insurer	N	$= 0$	$(0, 1)$	$= 1.0$	$(1, 1.5]$	> 1.5
All insurers	596,608	0.019	0.067	0.364	0.328	0.221
UnitedHealthcare	182,999	0.017	0.020	0.289	0.376	0.298
Cigna	98,765	0.009	0.029	0.148	0.311	0.503
Aetna	55,102	0.046	0.044	0.903	0.005	0.003
UMR	53,810	0.009	0.011	0.042	0.928	0.011
Anthem BCBS	23,838	0.034	0.088	0.768	0.082	0.028
Florida Blue	17,418	0.034	0.119	0.840	0.005	0.001
Anthem Virginia	17,128	0.033	0.058	0.578	0.266	0.066
BlueCross BlueShield of Tennessee	14,667	0.001	0.013	0.002	0.552	0.432
Provider's response: ln(*provider ask*) *on the insurer's protest share* q						
CPT × quarter fixed effects					0.319 (0.148)	
CPT × quarter and provider-group fixed effects					0.254 (0.099)	

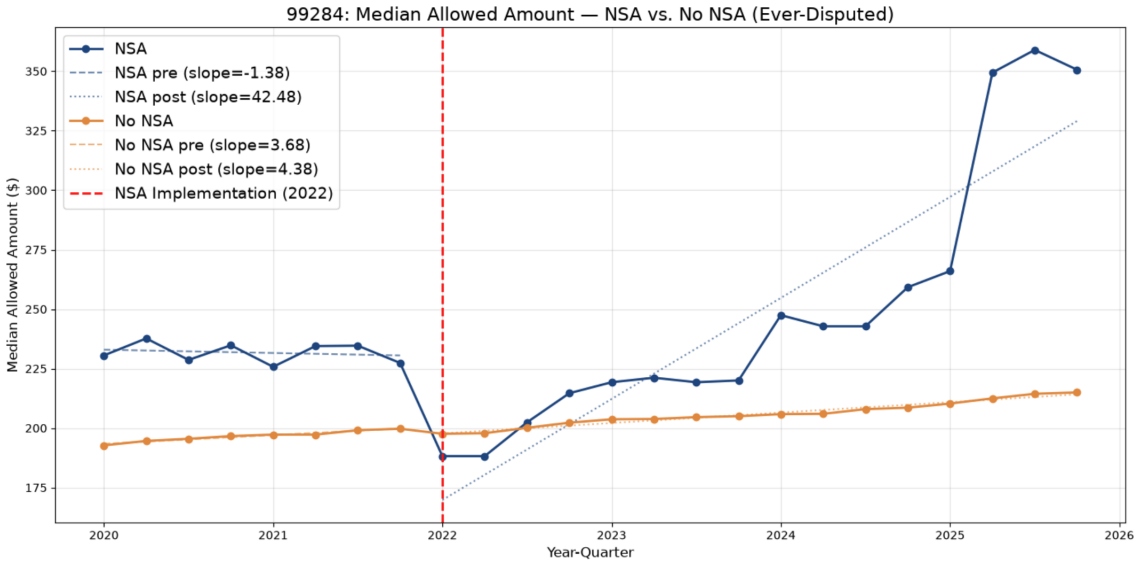
Notes: CMS public IDR files, 2023–2024; CPT 99283–99285 at place of service 23; disputes with both offers and an outcome; top 1 percent of the average offer trimmed ($N = 596{,}608$). Offers are expressed as multiples of the QPA. A *protest* is an insurer offer of exactly 0 or exactly 1.0 times the QPA; a *contest* is an offer above 1.0. Providers win 83.4 percent of all disputes, 85.6 percent when the insurer protests and 83.5 percent when it contests. The eight insurers listed account for 78 percent of disputes. Bottom panel: coefficient on q, the share of the counterparty insurer's *other* offers in the same CPT code and quarter that are protests (leave-one-out), with standard errors clustered by insurer (681 clusters; $N = 593{,}648$). Identification is from differences across insurers in protest propensity, not from whether the individual dispute was protested.

Figure 3: Binned Scatterplot of Share of Practice Revenue in Top 20 IDR Codes Versus Number of Cases



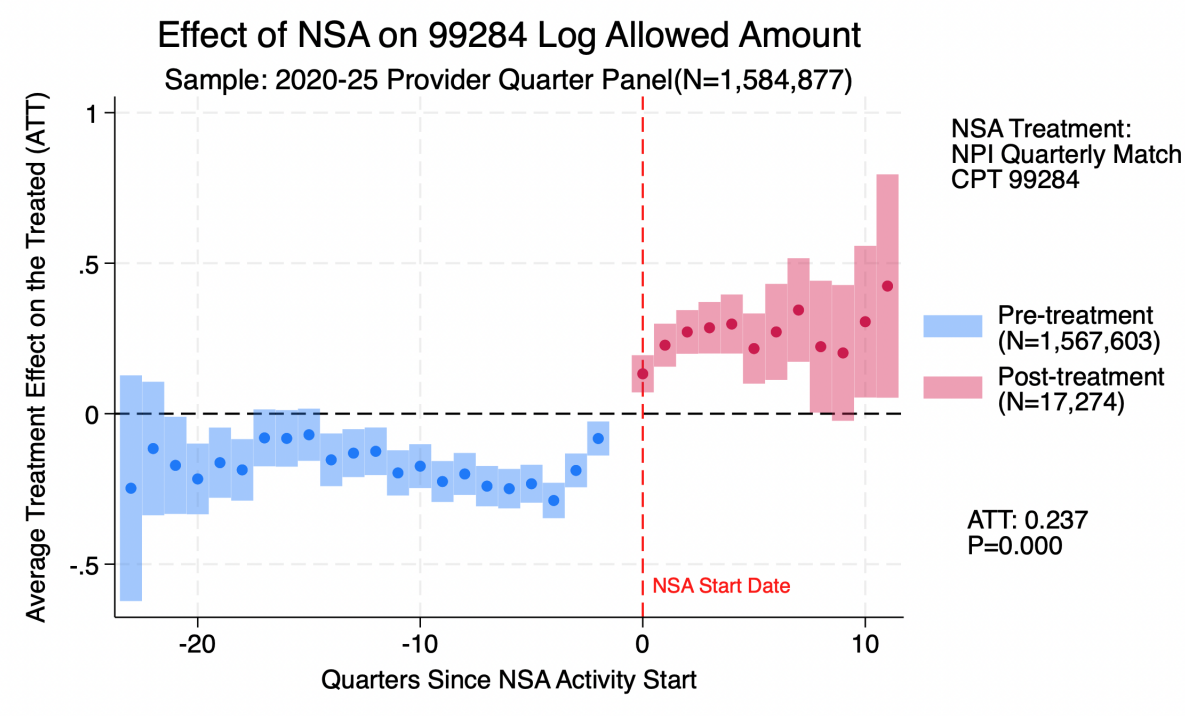
Note: Providers who are more exposed to the services most frequently disputed under the NSA are also the providers who bring the most cases. Figure 3 shows a binned scatterplot and the positive relationship between the percentage of a practice NPI's total revenue that comes from the top 20 CPT codes that are disputed under IDR and the number of cases they bring. We calculate the share of revenue coming from total allowed amounts and allowed amounts for the top 20 CPT codes using the Medicare Part B Carrier Data from 2022. While not causal, this result provides an interpretation that the providers who are more exposed to arbitration-related services have a higher propensity to engage with IDR.

Figure 4: Unadjusted Trends in In-Network ED Physician Prices Between NSA Participants and Non-Participants



NSA: 17,541 prov-qtrs, 1,029 providers (Pre: 4,769 | Post: 12,772)
No NSA: 1,677,226 prov-qtrs, 187,028 providers (Pre: 496,272 | Post: 1,180,954)

Figure 5: Differences in In-Network ED Physician Prices Following NSA Arbitration



8 Appendix

8.1 Background on Surprise Bills and Arbitration

So-called "surprise" bills occur when a patient receives medical services that are seemingly covered by insurance (e.g., in-network), but a component of the service is out-of-network and thus not covered by insurance. A number of widely-read media reports led to concerns in state legislatures that some providers and groups were intentionally not participating in insurer networks and instead exploiting surprise bills to increase revenues, and research suggested this was likely (Cooper, Scott Morton, and Shekita, 2020). Several private equity-backed provider groups had prominent reports written about them at this time.[14] These stories and concerns led to state action to limit patient exposure to surprise bills.

Initial surprise bill protections were state-led. To resolve payment disputes between insurers and providers, state enacted a range of polices. Formulaic payment rates (like those enacted in California) tie payment to an insurer's median in-network rate or other benchmark. Arbitration-based models emerged in New York and spread to other states like New Jersey, Texas, and Washington. Finally, a suggestion raised in Connecticut would simply make it illegal to send surprise bills to patients, but set up no payment standard, leaving hospitals and doctors to arrange contracts internally through side payments, a form of *ex-post* bargaining discussed favorably by Grossman and Hart, 1986. A key limitation of state policies is that they do not apply to self-funded employer plans.[15]

In response to the limitation that state laws did not protect self-funded plan beneficiaries[16] and to extend surprise bill protections to other states, Congress passed the No Surprises Act during the first Trump administration. The NSA has two important components. First, patients are no longer subject to surprise bills. The second major component is about determining insurer prices. Because these services are outside of an insurer's network, there is not a negotiated price between

[14] These include reports by the Congressional Research Service and the Government Accountability Office on "balance billing" and surprise bills. (e.g., https://www.congress.gov/crs-product/R46856 and https://www.gao.gov/assets/700/698332.pdf).

[15] https://www.healthaffairs.org/content/forefront/analyzing-new-bipartisan-federal-legislation-limiting-surprise-medical-bills

[16] In the US, self-funded plans are exempt from state regulations and are regulated by the federal government.

the two parties (providers and insurance companies). Naturally, providers wish for higher payment rates than insurers are willing to offer. If the two parties could agree upon rates, they would have agreed to a negotiated contract and the providers would be in-network.

To resolve these disputes, the NSA established an arbitration process modeled after exiting state policies. In NSA arbitration, either providers or insurers can bring a claim to resolve a balance bill dispute using "baseball-style" arbitration.[17] Importantly, NSA arbitration is "final offer" arbitration and parties are not able to adjust their bids. The arbitrator selects a single winning bid from either offer.

[17]This is frequently called "baseball style" arbitration as it is used in salary disputes between major league baseball players and the teams for whom they play.

8.2 Appendix Tables and Figures

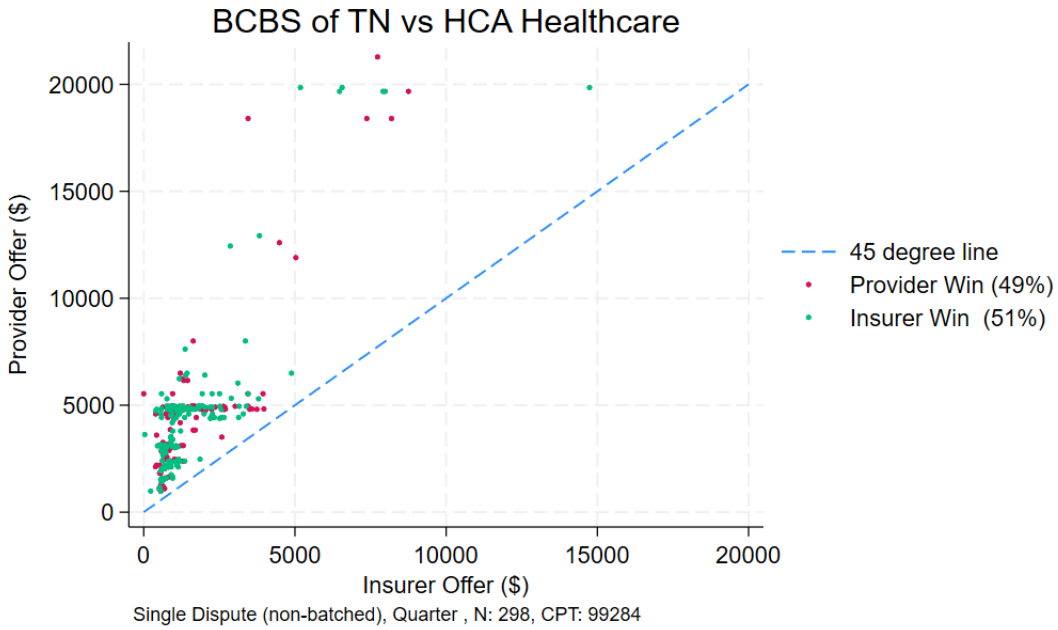
Table A1: Provider Win Rates and Provider Offers (relative to QPA)

Panel A: Outcomes		
Outcome	Cases	%
Insurer Win	88 641	15.82
Providers Win	471 835	84.18
Total	560 476	100.00

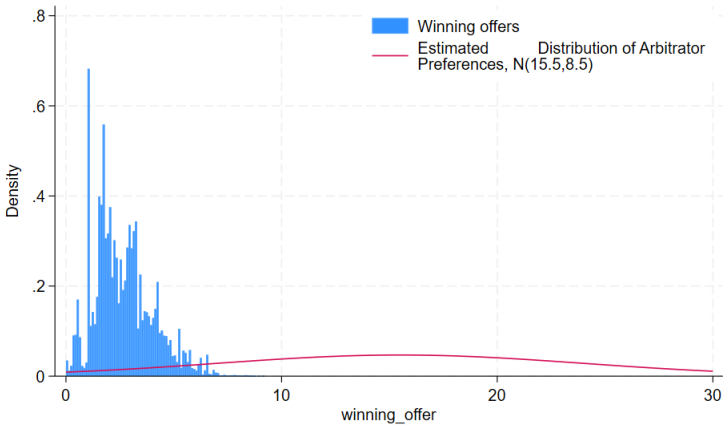
Panel B: Winning Provider Offers (×QPA)				
25th pct.	Median	75th pct.	Mean	SD
2.04	2.98	4.09	3.30	1.66

Note: $N = 471{,}835$. Provider offers are expressed as multiples of the Qualifying Payment Amount (QPA).

Note: Arbitration outcome data for this analytic sample is restricted to top 3 CPT codes (99283, 99284, 99285) which were performed in emergency room settings and is restricted to large providers and cases where provider offer is greater than insurer offer, and top 1% of average wage offers removed.

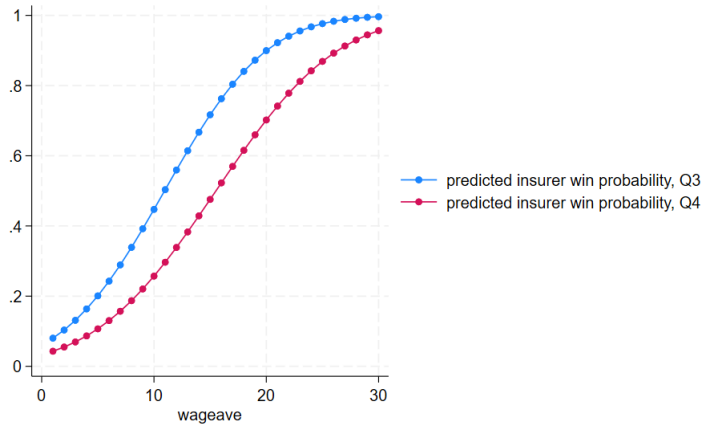


Appendix Figure 1: HCA vs BCBS of TN arbitration scatterplot



Appendix Figure 3: Estimated Density Function of Arbitrator Preferences

Note: Estimated distribution fit over actual winning offers. The probit parameters are estimated over a support mainly in the left tail. As we develop more fully with our inducement model, the traditional model is not only a poor fit but likely insufficient to capture party bidding behavior.



Appendix Figure 2: Estimated CDF of arbitrator preferences Q3 and Q4 of 2023

Note: Over time, it appears as if the arbitrator CDF is shifting leftward; this is consistent with a greater degree of inducement in later periods in the data.

8.3 Inducement Model Estimation

Section 3.3 presents a model of arbitrator inducement, in which arbitrators maximize a utility function that puts some weight α on the getting the "correct" outcome and the rest of the weight, $(1-\alpha)$, on induced revenue, where revenue is the arbitrator fee times the number of cases. When the IDRE maximizes this utility function, the optimal selected choice is: $\mu' = \mu + \frac{(1-\alpha)fk}{2\alpha}$. I.e., the arbitrator increases his chosen outcome away from the "correct" one and does so more when the fee per case is higher and/or when the ability to induce more cases is stronger. The second term is an upward shift of inducement, i.e., $\Delta = \frac{(1-\alpha)fk}{2\alpha}$. If one assumes the whole term above is what the arbitrator maximizes, you can estimate the inducement model with the probit.

Rearranging what's above (the arbitrator's ideal outcome) results in $\mu' - \mu = \frac{(1-\alpha)fk}{2\alpha} \leftrightarrow \alpha = \frac{fk}{2(\mu'-\mu)+fk}$, i.e., a solution for α. With variation in arbitrator fees, one can identify different alphas for each arbitrator. But a problem is that CMS doesn't release the identity of the arbitrator in the data set; only their fee. To address this limitation, we use the Internet Archive Wayback Machine and acquire 2023 fees for single and batched determinations for each arbitrator[18] and merge arbitrator identities with the public NSA data set. Fee matches aren't perfect; for example FHAS and National Medical Reviews both charge $490, so we can't link them to that fee. However, Keystone is the only IDRE to charge $600 for a single-determination in 2023, so we can clearly attribute these cases to Keystone. Matching where we can, and modifying the probit regression to include arbitrator fixed effects, allows us to recover an arbitrator-specific optimal mean μ' (at least for the IDREs that we can match).

Once we have the arbitrator-specific μ', we can back out an arbitrator-specific α, i.e., how much weight each puts on inducement. We are required to make a few assumptions to do this. First, we assume that absent any inducement, the optimal result would be the QPA, e.g., $\mu = 1$. Second, k is a free (floating) parameter, so we calibrate it at 0.8 (although it can be simulated at other values). It represents an elasticity of inducement, or how strongly

[18] https://web.archive.org/web/20230216150958/https:/www.cms.gov/nosurprises/help-resolve-payment-disputes/certified-idre-list

arbitrators can convert parties to use the NSA system. Another way to conceptualize k is as a competition vs. collusion conduct parameter—it's the same for all IDREs, so if they all are cooperating to keep the system going, then they can induce a lot of cases from providers and still retain high fees, but if there's lots of entry/competitive threat from other IDREs who are more charitable, it can undercut this incentive or elasticity of inducement. Finally, since the CMS data is not in dollars but is relative to the QPA, but the fee is in dollars, we standardize the fee to relative-to-QPA by dividing by $200, i.e., the QPA median for 99284 (and only estimate the model on 99284). We observe variation across arbitrators that shows: (A) the different arbitrators vary by a factor of nearly two from top-to-bottom in their α (when we've fixed k at some number) and also (B) that all the arbitrators have "low" alphas (where $\alpha = 1$ means no inducement and $\alpha = 0$ is exclusively inducement), given they all have low alphas.

8.4 Risk Aversion and Win Rates in Final-Offer Arbitration

This appendix derives the quantitative bound used in Section 3.2 showing that implausibly large risk aversion is required to generate the observed provider win rates in NSA arbitration.

8.4.1 Environment

Consider the final-offer arbitration game described in Section3.1. The insurer submits offer w_i and the provider submits offer w_p, with $w_p \geq w_i$. The arbitrator selects the offer closest to their ideal point w_a, which is drawn from a continuous distribution with CDF $F(\cdot)$ and density $f(\cdot)$.

The probability that the insurer's offer is selected is given by:

$$P(w_i, w_p) = F\left(\frac{w_i + w_p}{2}\right)$$

and the provider's offer is selected with probability $1 - P(w_i, w_p)$.

8.4.2 Party Preferences

We assume that arbitration parties maximize expected utility over final wealth, and not solely the arbitration award. Let b_i and b_p denote baseline wealth of the insurer and provider, respectively. Expected utilities are

$$EU_i(w_i, w_p) = PU_i(b_i - w_i) + (1 - P)U_i(b_i - w_p)$$

and

$$EU_p(w_i, w_p) = PU_p(b_p + w_i) + (1 - P)U_p(b_p + w_p)$$

CARA (Constant Absolute Risk Aversion) utility is given by:

$$U_j(x) = -exp(-\gamma_j x), \quad j \in \{i, p\}$$

where $\gamma > 0$ denotes absolute risk aversion.
Deriving the insurer FOC,

$$\frac{\partial EU_i}{\partial w_i} = \frac{1}{2} f\left(\frac{w_i + w_p}{2}\right)[U_i(b_i - w_i) - U(b_i - w_p)] - PU_i'(b_i - w_i) = 0$$

, and substituting CARA utility and canceling terms yields

$$\frac{1}{2} f\left(e^{\gamma w_i} - e^{\gamma w_p}\right) = P\gamma e^{\gamma w_i}$$

Similarly, differentiating EU_p with respect to w_p gives:

$$\frac{1}{2} f\left(e^{-\gamma w_i} - e^{-\gamma w_p}\right) = (1 - P)\gamma e^{-\gamma w_p}$$

These two FOCs jointly characterize equilibrium offers and win probabilities under CARA preferences.

8.4.3 Small-Stakes Approximation

Arbitration stakes under the No Surprises Act are small relative to baseline firm wealth. We therefore consider the approximation $\gamma|w_j| \ll 1$ which implies

$$e^{\gamma w} \approx 1 + \gamma w$$

Applying this to the first FOC gives:

$$\frac{1}{2} f\gamma \left(w_i - w_p\right) \approx P\gamma$$

Canceling γ gives

$$\frac{1}{2} f \left(w_i - w_p\right) \approx P$$

An analogous approximation of the other FOC gives

$$\frac{1}{2} f \left(w_i - w_p\right) \approx 1 - P$$

8.4.4 Implication for Win Probabilities

This implies that deviations of win probabilities from one-half depend on the product $\gamma\Delta w$ where $\Delta w \equiv w_p - w_i$

More generally, under smooth distributions of arbitrator ideal points, equilibrium win probabilities can be locally approximated by a logistic expression

$$P = \frac{1}{1 + exp(-\gamma\Delta w)}$$

Solving for the level of risk aversion required to generate a given win probability P yields:

$$\gamma\Delta w = log\left(\frac{P}{1-P}\right)$$

To generate a provider (or insurer) win probability of 80 percent, the above equation requires

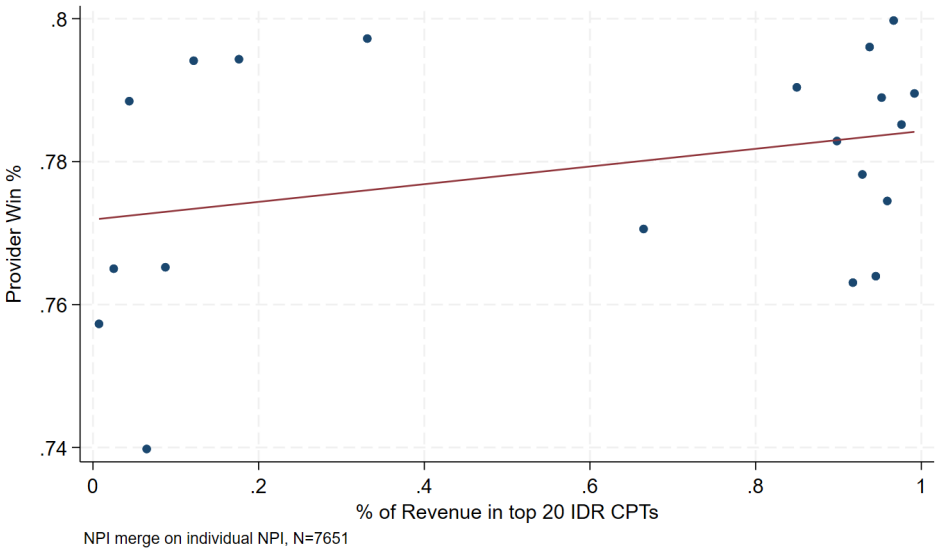
$$\gamma \Delta w = log\left(\frac{0.8}{0.2}\right) = log(4) \approx 1.386$$

In the NSA context, typical arbitration disputes involve Δw on the order of a few hundred dollars. Satisfying the approximation above therefore requires values of γ large enough that the small-stakes approximation underlying CARA preferences becomes implausible for large, diversified firms. This establishes the quantitative bound used in Section 3.2.

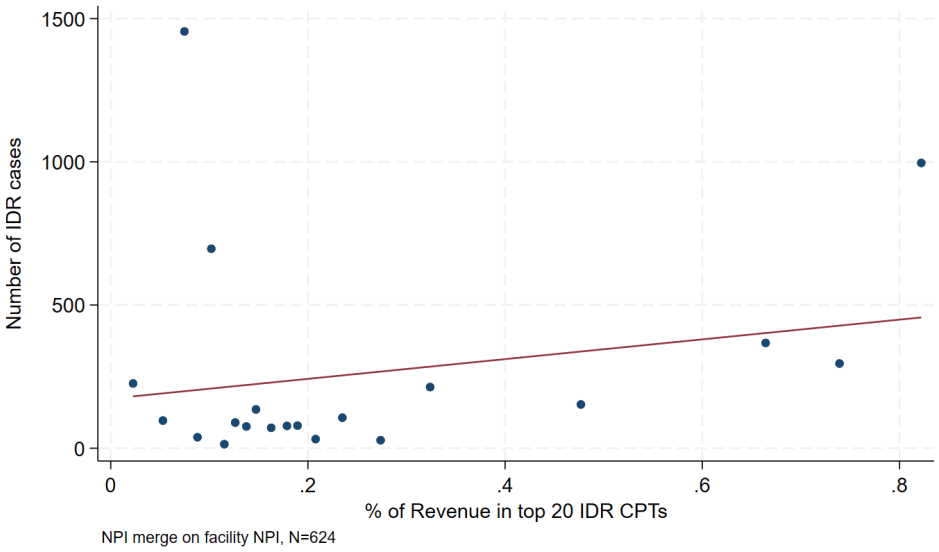
Under CARA utility, asymmetric win rates in final-offer arbitration depend on the interaction between risk aversion and the scale of arbitration stakes. Achieving win probabilities near those observed in the data requires $\gamma \Delta w \approx 1.4$ which is difficult to reconcile with the magnitude of NSA disputes relative to baseline firm wealth. This motivates the alternative mechanism developed in Section 3.3

8.5 Provider Exposure to NSA Services

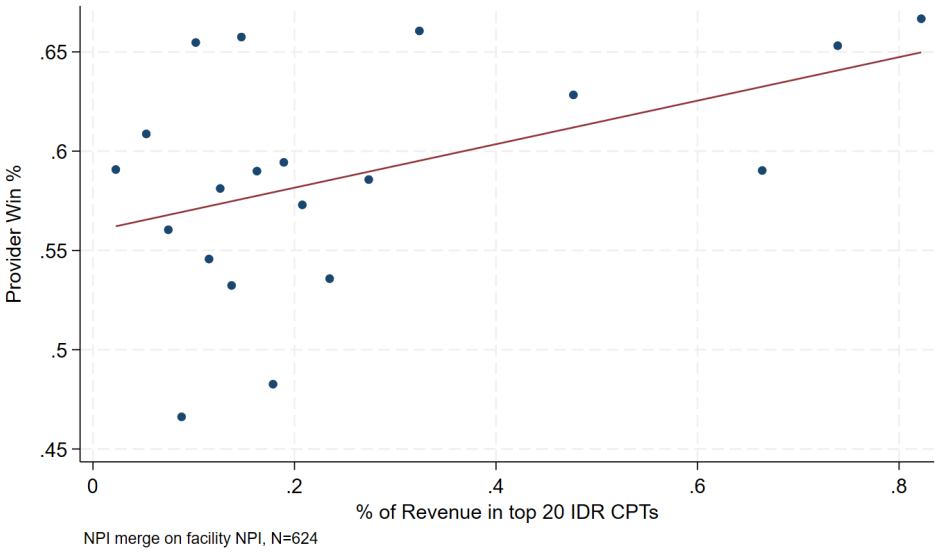
We build on Figure 3 for facility as well as individual NPIs, and for win rate shares as well as count of total cases. Finally, we also show that the win rate for HCA, a large hospital chain where most of their business is not related to surprise billing, the win rate is nearly 50%, as opposed to the lopsided provider win rates in the aggregate data. This evidence is suggestive that the providers most exposed to arbitration's effect on revenue are also engaging in the system most often.



Appendix Figure 4



Appendix Figure 5



Appendix Figure 6