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

Benjamin Chartock* *Bentley University*

Christopher Whaley† *Brown University*

May 8, 2026

Abstract

Hundreds of thousands of arbitration disputes between doctors and health insurers occur yearly under the No Surprises Act, impacting firm profits and affecting health care premiums through arbitration's impact on bargaining leverage. In this paper, we use public data to estimate a model of arbitration, find that insurers lose 84% of cases and investigate the sources of the lopsided arbitration results. We rule out differences in risk aversion that could drive observed outcomes, but cannot rule out an arbitrator incentive to induce more cases, a form of regulatory capture that policymakers can ameliorate, and which offers lessons on arbitration generally.

*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, 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 can have desirable properties in solving incomplete contracts (Williamson 1975), which is why it is observed in settings as widely variable as cell phone plans, baseball players' contracts, international trade disputes, and child custody battles.[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.[2] Yet economists know little about the underlying mechanics of arbitration and the factors that influence arbitrators' decisions. A key barrier to understanding arbitration is that, by nature, most arbitration processes are sealed and settlement terms are shielded from public view. Because comparatively little is known about the arbitration mechanism, it is hard to evaluate its effectiveness; this is the context of this paper.

We examine the mechanics of arbitration through the lens of the No Surprises Act (NSA), a federal law passed in 2020 and implemented in 2022, that adds consumer protections against "surprise" bills from medical providers. To resolve contractual disputes between providers and health insurers that underlie surprise bills, the NSA instituted what is likely the largest final-offer arbitration system ever instituted in the US. This system has resulted in over 3 million arbitration cases between US health care providers and health insurers in just the first three years. We leverage granular data on each of these disputes to examine how arbitration incentives operate in this important context.

To do so, we first describe a traditional economic model of arbitration, which has the dual purpose of highlighting the economic incentives of involved parties but also providing a lens for interpreting how this policy might affect negotiated private health insurance rates in the United States. We next tie this model to the data, estimating the parameters of the model with public data on 500,000+ arbitration cases, final offers, and results released by the Centers for Medicare and Medicaid Services (CMS). However, the results we recover from the canonical model fit poorly and imply unrealistic behaviors by both arbitrators and dispute participants.

[1] https://files.consumerfinance.gov/f/201503_cfpb_arbitration-study-report-to-congress-2015.pdf
[2] https://files.epi.org/pdf/144131.pdf

We investigate two potential mechanisms that can explain this inconsistency. The first relates to risk. In settings where parties diverge in their preferences for risk, asymmetric arbitration win rates may be observed (and which indeed which we find). However, we show asymmetric risk preferences are unlikely to explain observable patterns in the data. The second possibility we investigate is regulatory capture — namely, one party might place undue influence on the process itself. We propose a variation of the model which nests both the incentives for arbitrators to reach a decision consistent with the terms of the law, and also places some utility weight on revenue earned from additional cases resulting from regulatory capture. In short, we allow for arbitrators to endogenously induce cases through their decision outcomes. We then use financial data on ownership information from both arbitrators and the firms engaged in disputes with those same arbitrators, and show that there is vertical integration between arbitrators and the very firms they engage with in cases. We estimate higher inducement rates for these firms, which we can link to public CMS data by matching on fees and calendar year, something previously not done as best as we can observe.

For context on surprise medical bills and this arbitration system, these bills have been a frequent occurrence for patients and are defined as when a patient goes to an in-network hospital, is treated by an in-network primary surgeon, and yet an ancillary provider, such as radiologist or emergency room doctor treats them, but is not an in-network member of their insurance plan. The patient is then responsible for a higher out-of-network bill from the ancillary provider. The incidence and disruption caused by surprise billing prompted substantial policy action, first at the state level and then culminating in the NSA, which took effect in 2022. Government projections anticipated 20,000 arbitration cases per year, yet more than 3 million disputes have been filed.[3] In its first two years alone, the NSA generated roughly $5 billion in total costs (Hoadley and Watts 2025). Previous summaries of these cases have noted the large share of cases won by providers.

In our analysis of NSA arbitration data, we find that 84% percent of cases are won by providers. Additionally, providers are "winning" NSA arbitration disputes at approximately

[3]https://www.hfma.org/payment-reimbursement-and-managed-care/no-surprises-act-arbitration-has-been-a-bonanza-for-a-few-provider-groups/

3 times the qualified payment amount (QPA) that is designed to serve as a benchmark for arbitrators. While these statistics characterize how the NSA market has initially operated, the underlying mechanics behind the large discrepancy between win rates and outcomes is not well understood. Accordingly, we do not rule out behavior consistent with regulatory capture, specifically the inducement of a high number of arbitration cases which are financial beneficial to involved parties. Consistent with our model of regulatory capture, we also note that many IDRE entities have financial ties to private equity (PE) companies, many of which also sponsor or own provider groups that bring NSA cases.

This paper relates to a number of literatures, but especially a growing literature on the corporatization of healthcare (Atul Gupta, Howell, Yannelis, and Abhinav Gupta 2024; Lin, Munnich, Richards, C. Whaley, and Zhao 2022; Gandhi and Olenski 2024; Atul Gupta, La Forgia, and Sacarny 2024; Diwan, Eliason, League, Leder-Luis, McDevitt, and Roberts 2025; Singh, Radhakrishnan, Adler, and C. Whaley 2025). We contribute to the literature on how networks form and how rates are set via bargaining (Gowrisankaran, Nevo, and Town 2015; Kate Ho and Lee 2017; Barwick, Swanson, and Xia 2025;Prager and Tilipman 2025). In particular, we are most closely related to the Prager and Tilipman paper as we provide a microfoundation for non-zero disagreement payoffs. We contribute to a growing literature on the intersection of health care, labor markets, and rising use of artificial intelligence, as the surprise billing providers such as radiologists operate business models severely impacted by the growth of AI (Dranove and Garthwaite 2022; Gottlieb, Polyakova, Rinz, Shiplett, and Udalova 2023; Acemoglu, Autor, Hazell, and Restrepo 2022). And finally, we reawaken an old labor and law & economics literature about final-offer arbitration (Farber 1980; Ashenfelter and Bloom 1984; Gibbons 1988) that has been understudied for many years as union membership has fallen.[4].

Although we refrain from normative policy suggestions beyond proposing greater use of Transparency-in-Coverage data by arbitration entitites and stricter limits on shared corporate structure between arbitrators and parties, our back-of-the envelope calculations suggest

[4]The exception is a recent paper which extends Gibbons' model to empirical tractability relying on the same data source studied in the older series of arbitration papers, yet adapted to include Bayesian learning by the arbitrator and empirical estimation (Kong, Silveira, and Tang 2025)

a potential $2.7 billion increase in provider payments for just a single emergency medicine procedure.

Our paper proceeds as follows. Section II presents a model of arbitration, Section III discusses the data, Section IV shows the results, and Section V concludes.

2 Model

This section develops a sequence of models to study outcomes in No Surprises Act final-offer arbitration. We begin with a standard benchmark model that delivers symmetric equilibrium outcomes and equal win rates. We then extend the model to allow bargaining parties to be risk averse over final (post-arbitration) wealth, following prior work, and show that while risk aversion can generate asymmetric win rates in principal, it does so only under implausible parameter values given the scale of NSA disputes. We therefore introduce a model in which arbitrator incentives affect outcomes. This inducement model provides a parsimonious mechanism capable of generating the large and persistent provider win rates observed in the data and serves as the basis for our empirical analysis.

2.1 Baseline Model of Final-Offer Arbitration

We begin by introducing a model of 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.[5] In the model, insures and providers offer w_i and w_p, respectively, to the arbitrator.[6]

[5] Notationally, we consider this bargaining over a wage as doctors get paid by insurers, and most if not all other arbitration papers study wage disputes such as what occurs between municipalities and fire and police departments (where arbitration is used because fire and police departments are legally forbidden from using strikes)

[6] In theory, both parties can initiate an arbitration, although in practice, the providers initiate more frequently. This will become relevant later when we introduce arbitrator inducement.

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, w, based on his preferred outcome w_a. The arbitrator's utility function is modeled as a quadratic loss from the distance between their 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 $w \in [w_i, w_p]$ closer to w_a; in other words, the arbitrator will pick the insurer's offer if w_a is less than the average of the offers, $w_a \leq (w_i + w_p)/2$, or if w_a is less than halfway between the 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 the parties only know the arbitrator's preferences up to a distribution. Parties understand this uncertainty as a random variable with cumulative distribution function $F(w_a)$ and probability density function $f(w_a)$. The probability that the insurer's offer is chosen is $F(w_a)$, and because the arbitrator must choose either one or the other, the probability of the provider being chosen is simply $1 - F(w_a)$.

$$Pr(insurer\ win) = Pr(w_a \leq \frac{w_i + w_p}{2}) = F(w_a) \tag{1}$$

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:

$$\min_{w_i} \quad w_i \cdot F\left(\frac{w_i + w_p^*}{2}\right) + w_p^* \cdot \left[1 - F\left(\frac{w_i + w_p^*}{2}\right)\right] \tag{2}$$

$$\max_{w_p} \quad w_i^* \cdot F\left(\frac{w_i^* + w_p}{2}\right) + w_p \cdot \left[1 - F\left(\frac{w_i^* + w_p}{2}\right)\right] \tag{3}$$

Solving via first order conditions:

$$F\left(\frac{w_i^* + w_p^*}{2}\right) = \frac{1}{2} \tag{4}$$

$$\frac{1}{f\left(\frac{w_i^* + w_p^*}{2}\right)} = (w_p^* - w_i^*) \tag{5}$$

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 (that is, the reciprocal of the density function is larger). 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. In the arbitration context, a provider will trade off offering slightly higher (and earning more if they win) with reducing their probability that their offer is selected. The insurer has a symmetric and opposite incentive to bid low but not so low that the arbitrator thinks they are "being greedy" or demanding too much and the other party prevails.

2.1.4 Estimation

For estimation, we follow existing literature and assume the arbitrator's preferred settlement is a random variable distributed $x \sim N(\mu, \sigma)$. Nash Equilibrium offers are then:

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

As first shown by Ashenfelter and Bloom 1984, the parameters that govern the distribution of arbitrators' preferences can be recovered from a probit regression corresponding to Equation (1), where the outcome variable is a binary dependent indicator of whether the insurer won. The coefficient on the regressor of average wage offer, $(w_i + w_p)/2$, provides an estimate of $1/\sigma$ and the coefficient on the constant term provides an estimate of μ/σ, with publicly available facts of the case also potentially shifting the outcome. In our empirical work, we control for a variety of case-specific fixed factors like location of the dispute and insurer and provider fixed effects, which is consistent in other papers that study arbitration (e.g.,

Ashenfelter, Bloom, and G. Dahl 2013). We estimate the probit model using maximum likelihood.

Some intuition for our estimation procedure is as follows. Suppose that the insurer offers $100 and the provider offers $120, and the provider is chosen by the arbitrator. In another case, the insurer offers the same $100, but the provider instead raises his offer by a dollar to $121. When he does this, suppose the provider then loses. This provides information that somewhere between $120 and $121, the arbitrator believes that the provider has asked for too high a ruling, and the arbitrator switches from awarding in favor of the provider to awarding against them. With data on a series of final offers and an indicator of who was selected, we can trace out the distribution of the arbitrator's preferences with the help of assuming a functional form.

2.2 Risk Aversion over Final Wealth

A natural extension of the baseline model allows parties to be risk averse over final (post-arbitration) wealth rather than over the arbitration award itself. 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). In this subsection, we incorporate risk aversion into the model and show that, given the scale of disputes under the No Surprises Act, this mechanism likely doesn't account for the large provider win rates observed in the data.

2.2.1 Utility Over Final Wealth

We modify the baseline framework by allowing each party to maximize expected utility over total wealth following arbitration. Let b_i and b_p denote the insurer's and provider's baseline wealth, respectively. Given arbitration offers w_i and w_p, expected utilities are

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

and

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

Transfers are zero-sum, but unlike the baseline model, risk aversion now operates through total wealth rather than the arbitration award alone. As in the benchmark case, each party chooses its offer taking the other party's offer as given.

2.2.2 CARA approximation

To assess the quantitative importance of risk aversion, we approximate preferences using constant absolute risk aversion (CARA) utility,

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

where γ_j denotes the coefficient of absolute risk aversion. Under CARA preferences and standard assumptions on the distribution of arbitrator ideal points, equilibrium win probabilities depend on the product of risk aversion and the scale of arbitration stakes. Intuitively, risk aversion affects behavior only to the extent that changes in the arbitration outcome meaningfully alter total wealth.

2.2.3 Quantitative Bound

A simple quantitative bound illustrates the limited explanatory power of this mechanism. Solving the equilibrium conditions under CARA utility implies that achieving a win probability of 80 percent for one party requires $\gamma \Delta w \approx log(4)$, where γ is the relevant coefficient of absolute risk aversion and Δw is the equilibrium offer gap. In the NSA setting, typical arbitration stakes are on the order of a few hundred dollars per dispute. Satisfying this condition with such stake sizes implies values of γ that are orders of magnitude larger than those commonly estimated in labor, insurance, or industrial organization contexts.[7]

This implication is further weakened by institutional features of the NSA. Arbitration outcomes represent a negligible share of baseline firm wealth for both insurers and large

[7] Cohen and Einav 2007 estimate CARA utility in auto insurance deductible choices and find that most individuals are approximately risk neutral for gambles on the order of $100–$500, implying absolute risk aversion on the order of 10^{-4} to 10^{-3} per dollar for typical agents. Furthermore, Einav, Finkelstein, Ryan, Schrimpf, and Cullen 2013 study health insurance choice and find that observed behavior over price variation in the $100–$1,000 range implies limited curvature of utility and small risk premia, consistent with near risk neutrality for moderate stakes.

provider organizations who are diversified across many disputes. As a result, even substantial curvature in utility functions would induce only small behavioral responses in equilibrium offers and win probabilities.

2.2.4 Implication

Taken together, these considerations suggest that asymmetric risk aversion over final wealth cannot plausibly account for the large and persistent provider win rates observed in NSA arbitration. Although risk aversion can generate asymmetries in theory, doing so in this setting would require implausibly large risk-aversion parameters relative to the scale of the underlying disputes. We therefore turn to an alternative mechanism that operates through arbitrator incentives rather than party preferences.

2.3 Arbitrator Inducement and Regulatory Capture

The inability of risk aversion to account for the large provider win rates observed in NSA arbitration suggests that the source of asymmetry may lie outside the preferences of the disputing parties. We therefore consider a mechanism that operates through arbitrator incentives. Under the No Surprises Act, arbitration entities are compensated on a per-case basis and earn revenue per disputes. This creates incentives for arbitrators to consider not only the outcome of the current dispute, but also the impact of their decisions on future case volume.

2.3.1 Arbitrator Objective With Inducement

We model the arbitrator as maximizing a weighted objective that trades off selecting an outcome consistent with the arbitrator's ideal payment against revenue generated through induced disputes. Let x denote the arbitrator's ideal outcome absent inducement, and let w denote the selected payment. Arbitrator utility is given by

$$U_a = -\alpha(x - w)^2 + (1 - \alpha)fn, \qquad \text{where } n = k(w - x).$$

Here, k parametrizes the strength of the inducement effect – as the arbitrator selects wages higher than what is ideal, more providers take advantage of that wedge and submit more cases, thus increasing both n and their revenue, which is just count n times the fee f.

The parameter $\alpha \in [0, 1]$ governs the weight placed on selecting a "correct" outcome versus induced revenue. When $\alpha = 1$, arbitrators are purely outcome-oriented and the model collapses to the baseline case.

2.3.2 Optimal Arbitrator Behavior

Solving the arbitrator's problem yields an optimal selected payment

$$w^* = x + \frac{(1-\alpha)fk}{2\alpha}$$

When the IDRE maximizes this utility function, the optimal selected choice is as shown above, 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. Even modest inducement incentives $(1 - \alpha)$ therefore shift the arbitrator's preferred outcome upward relative to x. This shift is increasing in both arbitration fees and the effectiveness of inducement. Importantly, unlike risk aversion, this mechanism does not depend on the scale of the arbitration stakes relative to baseline wealth.

2.3.3 Mapping Inducement Into Win Rates

In a final-offer arbitration setting, shifts in the arbitrator's preferred outcome translate directly into changes in win probabilities. Because the arbitrator selects the offer closest to w^*, an upward shift in w^* increases the likelihood that the provider's offer is selected. When arbitrator ideal points are dispersed, relatively small shifts in w^* can generate large changes in observed win rates. In contrast to risk-aversion explanations—which require implausibly large curvature to generate 80–85 percent win rates—this inducement mechanism can generate substantial asymmetries with economically reasonable parameter values.

The inducement model yields testable implications that we examine in the data. Arbi-

trators with greater inducement capacity—due to experience, market position, or vertical integration—should exhibit systematically higher provider win rates. In the empirical analysis, we evaluate these implications using variation across arbitrators and fee schedules.

3 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.[8] 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

[8] 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

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).

3.1 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 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).

3.2 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 the median award is higher than only 1.5% of all negotiated rates. 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.[9] 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.[10] 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 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 frequently insurers bid zero or one times the QPA. This behavior could be rationalized purely by the baseline model—for example, if the insurer thought the arbitrator had preferences of 2-times QPA and the provider bid anything more than 4-times QPA, it is rational for the insurer to bid zero. However, as we discuss in the concluding section, we believe it is more likely that insurers are protesting what they perceive are problems with the implementation of the law; lawsuits about non-payment proliferate.[11]

4 Results

In this section we show results from our baseline structural model and sturtural 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, the first quarter where arbitration regularly occurred with no major administrative pauses,

[9] 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.

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

[11] https://litigationtracker.law.georgetown.edu/issues/no-surprises-act/

and finds a mean of 10.7 and a standard deviation of 9.7 times the QPA (Table 1, column 1). [12]

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. Rather than concluding the seemingly unlikely case that arbitrator's preferences for an ideal outcome of arbitration are actually ten times as high as the benchmark negotiated rate (QPA), we instead explore and rule out some alternative hypotheses of this lopsided win rate coupled with high arbitration awards through our model of arbitrator inducement.

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 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 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. [13]

In Table 2, we find meaningful variation in our α estimates, signaling variation in IDRE inducement.[14] For example, the "Island Peer Review" IDRE has the highest α estimate,

[12] 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.

[13] 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%.

[14] Our model is identified ordinally, not cardinally, as the parameter k is a floating parameter which we set equal to 1 but shifts around α.

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.[15] 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 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. Yet expanding this model to include the incentives of arbitrators themselves is more consistent with the observed outcomes.

These arbitrator incentives, which in our model contribute to the skewed win rates and award amounts for providers, 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 85 percent and a payoff

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

of approximately 3-times QPA, and thus an expected payoff of 2.7 times 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.[16] 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 year prior to the NSAs passage. This suggests that use of price transparency data can lower the incidence of costly arbitration disputes.

[16]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

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 regulatory problem, not a mechanism design flaw.

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

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo (2022). "Artificial intelligence and jobs: Evidence from online vacancies". In: *Journal of Labor Economics* 40.S1, S293–S340.

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.

Barwick, Panle Jia, Ashley T Swanson, and Tianli Xia (2025). *A Double Dose of Reform: Insurance and Centralized Negotiation in Drug Markets*. Tech. rep. National Bureau of Economic Research.

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.

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.

Diwan, Renuka, PJ Eliason, RJ League, Jetson Leder-Luis, Ryan C McDevitt, and James W Roberts (2025). "Competition and fraud in health care". In: *Work in Progress*.

Dranove, David and Craig L Garthwaite (2022). *Artificial intelligence, the evolution of the healthcare value chain, and the future of the physician*. Tech. rep. National Bureau of Economic Research.

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.

Gandhi, Ashvin and Andrew Olenski (2024). *Tunneling and hidden profits in health care*. Tech. rep. National Bureau of Economic Research.

Gibbons, Robert (1988). "Learning in Equilibrium Models of Arbitration". In: *The American Economic Review* 78.5, pp. 896–912. ISSN: 00028282. URL: http://www.jstor.org/stable/1807156.

Gottlieb, Joshua D, Maria Polyakova, Kevin Rinz, Hugh Shiplett, and Victoria Udalova (2023). *Who values human capitalists' human capital? The earnings and labor supply of US physicians*. Tech. rep. National Bureau of Economic Research.

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 (2017). "Insurer competition in health care markets". In: *Econometrica* 85.2, pp. 379–417.

— (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.

Lin, Haizhen, Elizabeth L Munnich, Michael R Richards, Christopher Whaley, and Xiaoxi Zhao (2022). *Outside Equity and Healthcare Firm Behavior*. RAND.

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 (Apr. 2025). *Regulating Out-of-Network Hospital Payments: Disagreement Payoffs, Negotiated Prices, and Access*. Working Paper 33727. Na-

tional Bureau of Economic Research. DOI: 10.3386/w33727. URL: http://www.nber.org/papers/w33727.

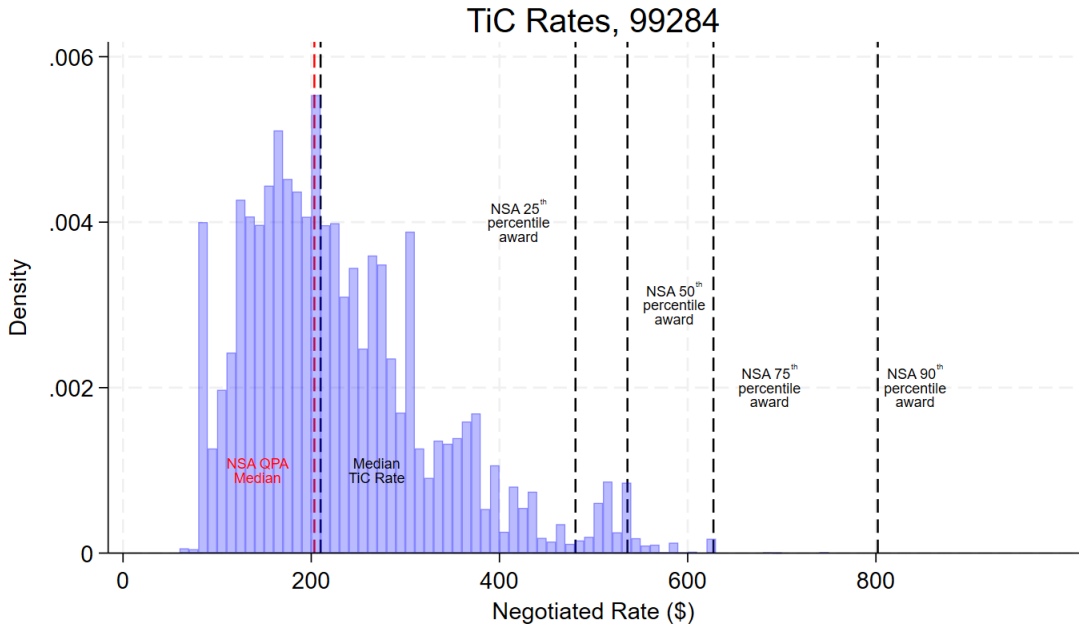
Singh, Yashaswini, Nandita Radhakrishnan, Loren Adler, and Christopher Whaley (2025). "Growth of private equity and hospital consolidation in primary care and price implications". In: *JAMA Health Forum*. Vol. 6. 1. American Medical Association, e244935–e244935.

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*.

6 Tables and Figures

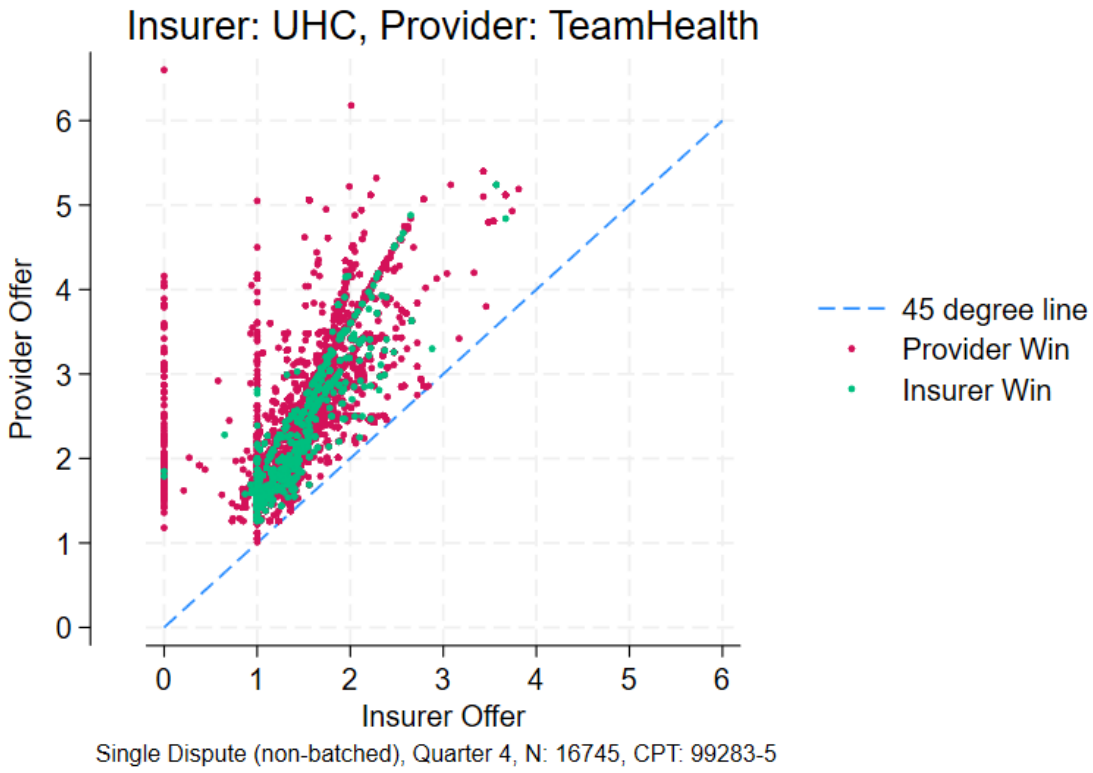
Figure 1: Arbitration Awards Relative to the QPA for Most Common Service

(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.

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.

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

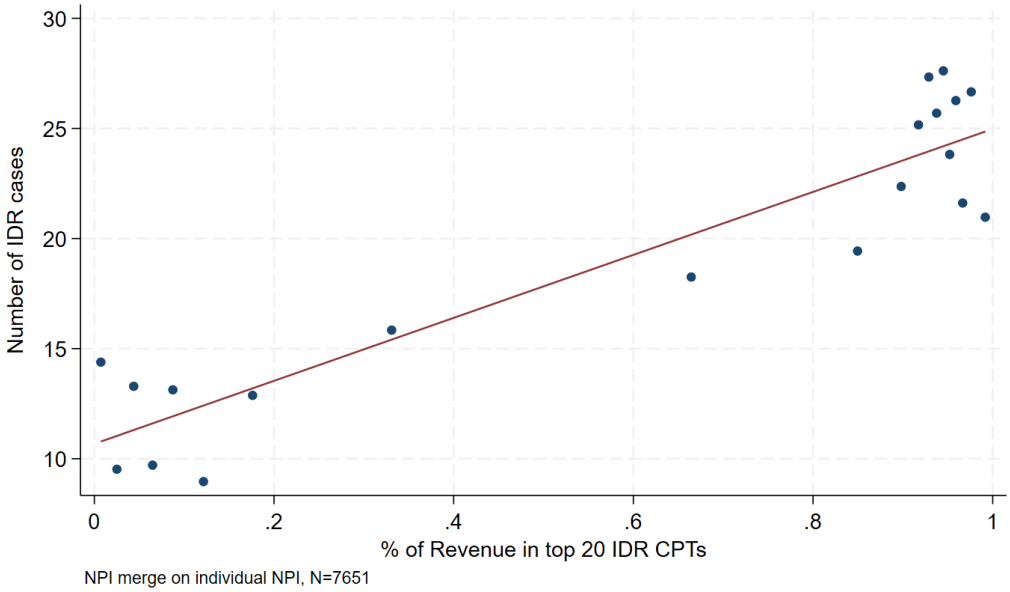
Note: Presented here are estimates from the structural model described in Section II. Quarters 1 & 2 were not retained due to starts & stops in the law subject to litigation brought by the Texas Medical Association in federal court.

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.

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.

7 Appendix

7.1 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.[17] 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.[18]

In response to the limitation that state laws did not protect self-funded plan beneficiaries[19] 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 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.[20] Importantly, NSA

[17] 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).

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

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

[20] 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.

arbitration is "final offer" arbitration and parties are not able to adjust their bids. The arbitrator selects a single winning bid from either offer.

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

Panel A: Outcomes

Outcome	Cases	%
Insurer Win	88641	15.82
Providers Win	471835	84.18
Total	560476	100.00

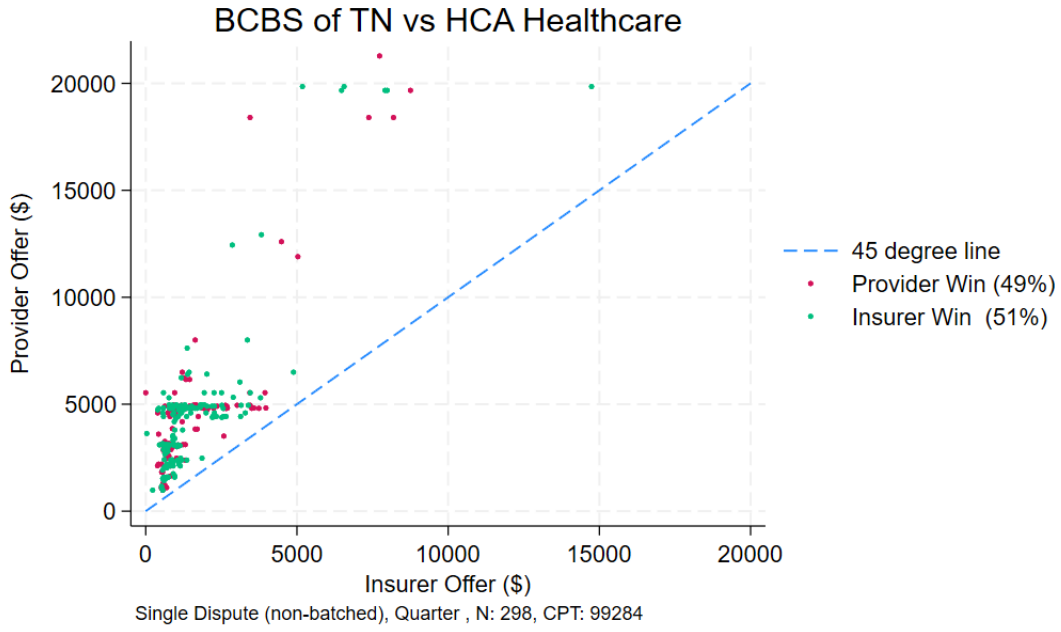
Panel B: Winning Provider Offers (×QPA)

25th pct.	Median	75th pct.	Mean	SD
2.04	2.98	4.09	3.30	1.66

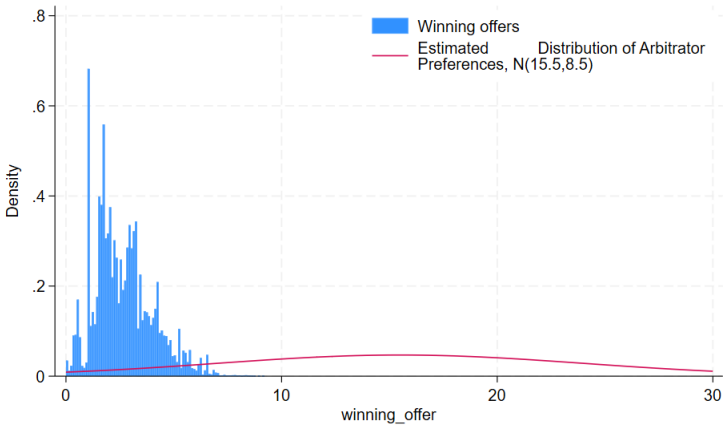
[flushleft]

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.

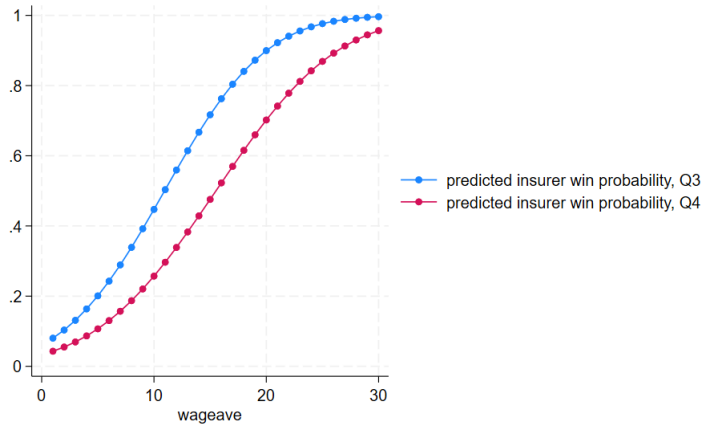


Appendix Figure 1: HCA vs BCBS of TN arbitration scatterplot



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.

7.2 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[21] 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

[21] 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.

7.3 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.

7.3.1 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 $1 - P(w_i, w_p)$.

7.3.2 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.

7.3.3 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

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

7.3.4 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

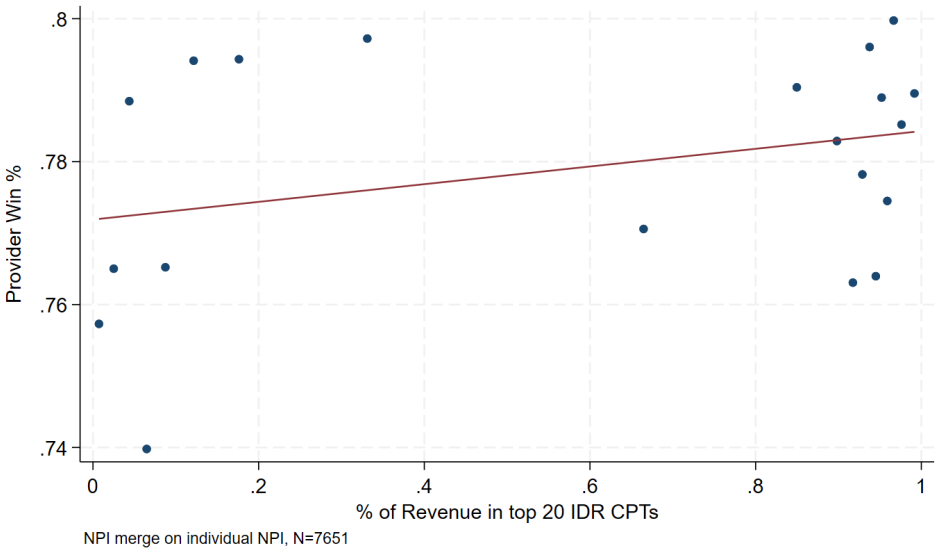
$$\gamma \Delta w = log\left(\frac{0.8}{0.2}\right) = log(4) \approx 1.386$$

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.

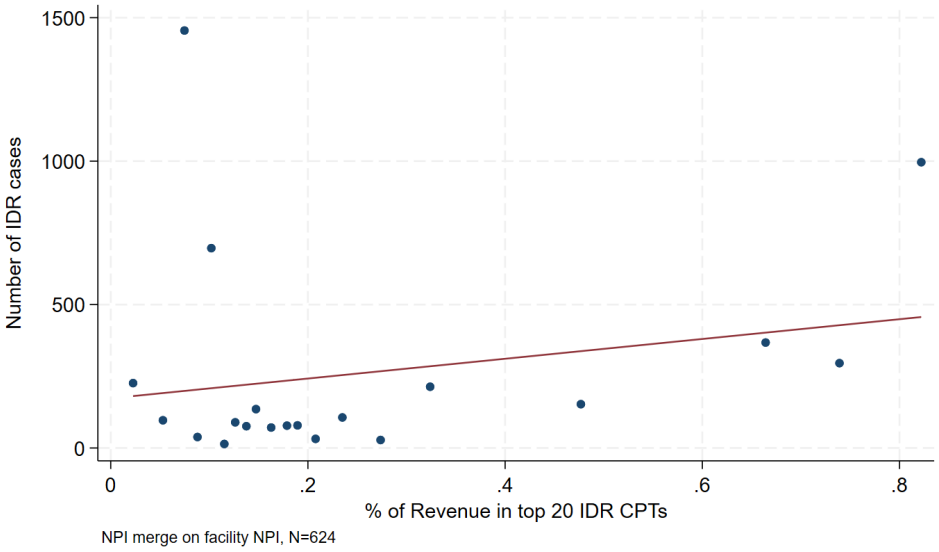
Under CARA utility, asymmetric win rates in final-offer arbitration depend on the interaction between risk aversion and the scale of arbitration stakes. Achieving win probabilities near those observed in the data requires $\gamma \Delta w \approx 1.4$ which is difficult to reconcile with the magnitude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

7.4 Provider Exposure to NSA Services

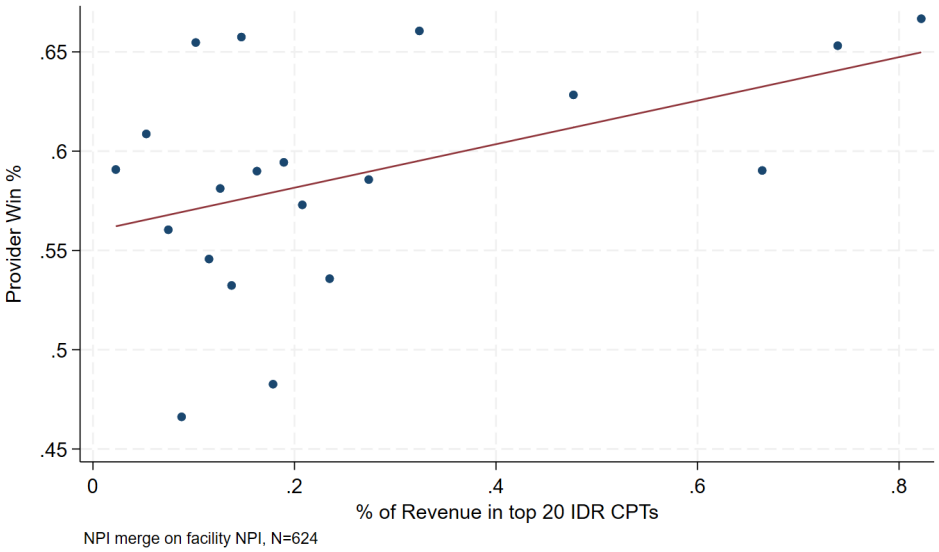
We build on Figure 3 for facility as well as individual NPIs, and for win rate shares as well as count of total cases. Finally, we also show that the win rate for HCA, a large hospital chain where most of their business is not related to surprise billing, the win rate is nearly 50%, as opposed to the 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