

RISK AND INFORMATION IN DISPUTE RESOLUTION: AN EMPIRICAL STUDY OF ARBITRATION

YUNMI KONG, BERNARDO S. SILVEIRA AND XUN TANG

ABSTRACT. We develop and estimate a structural model of arbitration, accounting for asymmetric risk attitudes and learning. Using data on public sector wage disputes in New Jersey, we compare the efficiency of two popular arbitration formats, final-offer (FOA) and conventional (CA). We find that, although CA hinders the transmission of case-relevant information from the disputants to the arbitrator, this format outperforms FOA by affording discretion to select awards. We also assess how risk-attitude differences between the disputants affect imbalances in arbitration outcomes, finding that risk aversion weakens a party's position in the dispute despite making them more likely to win arbitration.

Keywords: Arbitration, Dispute Resolution, Strategic Communication, Cheap-Talk, Risk Attitudes, Bargaining

JEL Classification: C57, C7, D82, D83, J52, K41

Date: May 30, 2024.

Kong (email: yunmi.kong.01@gmail.com) and Tang (email: xun.tang@rice.edu): Rice University; Silveira (email: silveira@econ.ucla.edu): University of California, Los Angeles. We are grateful to Yujung Hwang, Tasos Kalandrakis, Alessandro Lizzeri, Adriana Lleras-Muney, Maurizio Mazzocco, Isabelle Perrigne and Quang Vuong for helpful comments and suggestions. We would also like to thank seminar and conference participants at the Bank of Canada; Caltech; Carnegie Mellon; Chicago; Columbia; Cornell; KU Leuven; LSE; PUC-Rio; QMUL; Stanford; Toronto; UBC; UCLA; University of Melbourne; WVU; ASSA; Barcelona GSE Summer Forum, Applied IO; Bargaining: Experiments, Empirics, and Theory Workshop; Brazilian Econometrics Society Applied Economics Seminar; Cowles Conference on Models and Measurement; EEA-ESEM Virtual Congress; Empirical Models of Political Economy Conference; IAAE; IOOC; Interactive Online IO Seminar; Korean-American Economic Association Virtual Seminar; SEA; and SITE Session on Empirical Implementation of Theoretical Models of Strategic Interaction and Dynamic Behavior. Mary Beth Hennessy-Shotter at NJ PERC and arbitrators Ira Cure and Brian Kronick provided valuable information on police and fire arbitration practices in New Jersey. Special thanks to Ranie Lin, Shosuke Noguchi, and Jennifer Zhang for excellent research assistance. Sandy He, Susie Proo, Valeria Rojas, Heewon Song, Jinah Weon and Esther Yu contributed with the data collection.

1. INTRODUCTION

Arbitration is a private bilateral conflict resolution procedure in which a third party, the arbitrator, makes a binding decision on the dispute. Compared with formal litigation through a court system, arbitration is typically cheaper, faster and less formal. Moreover, arbitrators tend to be experts on the subject matter of the dispute, whereas judges assigned to court cases are usually generalists (Mnookin, 1998). Due to these advantages, arbitration has been extensively employed in the resolution of a variety of disputes including labor impasses, disagreements concerning commercial contracts, tort cases and tariff negotiations, among many others. In fact, Lipsky and Seeber (1998) surveyed the general counsels of the Fortune 1,000 companies in 1997, and found that 80 percent of the respondents had used arbitration at least once in the previous three years. Considering business-to-business disputes alone, the American Arbitration Association reported 9,196 cases in 2021, totaling over 15 billion dollars worth of claims.¹ In the public sector, as of the year 2000, around 30 states in the U.S. specified binding arbitration as the last-resort step in labor disputes for at least some categories of public employees (Slater, 2013).

There is substantial variation in arbitration formats, with two alternative designs—conventional and final-offer—standing out.² In each of these designs, the disputing parties submit to the arbitrator one offer each. The chief distinction is that in conventional arbitration the arbitrator is free to impose a ruling that differs from both offers, whereas in final-offer arbitration the arbitrator must select the offer of one side or the other. What is the relative performance of these two designs from a normative standpoint? A key dimension in such a comparison is the extent to which the arbitrator is able to acquire and use any pertinent information about the case at hand to deliver an appropriate ruling. In that sense, the choice between conventional and final-offer arbitration boils down to a trade-off between discretion and information transmission. On the one hand, conventional arbitration affords much more discretion to the arbitrator in making a decision, given the information that she has about

¹To be sure, these figures refer to actual disputes; the aggregate value of contracts that contain arbitration clauses is likely to be orders of magnitude larger.

²Conventional arbitration is the dominant format in consumer, commercial, and international arbitration, among others. Meanwhile, in addition to being the method of choice in salary disputes in Major League Baseball, final-offer arbitration has been employed by antitrust authorities to resolve disputes concerning high-profile merger cases, such as News Corp-DirectTV and the purchase of Adelphia by Time Warner and Comcast (Pecorino et al., 2021). In the setting that we analyze in this paper, public sector wage disputes, conventional and final-offer arbitration are arguably equally popular. Among the states using arbitration for this purpose, as of 2013, at least 14 employed final-offer arbitration (Carrell and Bales, 2013).

the case. On the other hand, final offer arbitration has the potential to facilitate the signaling of any private information the parties might have about the case through their offers, allowing the arbitrator to make a better-informed ruling. The reason is that offers in conventional arbitration are cheap-talk, whereas in final-offer arbitration they are not (Feuille, 1975; Gibbons, 1988).³ The cheap-talk nature of conventional arbitration incentivizes parties to make overly ambitious offers, which might reduce their informational content.⁴ Which of the two systems allows for better decisions is ultimately an empirical question.

This paper combines theory and empirics to compare the efficiency of the arbitrator’s decisions under the conventional and final-offer designs. To this end, we develop a new framework for the structural analysis of arbitration, employing data on wage arbitration between local governments and police and fire officer unions in the State of New Jersey. In this context, we define efficiency as the ability to deliver arbitration awards that are closer to the ideal or fair wage as interpreted by New Jersey law.⁵ We leverage our structural model and a transition of the default arbitration method in New Jersey from final-offer to conventional to measure the transmission of information under each arbitration format. Our results indicate that the information communicated in final-offer arbitration is more than twice as precise as that in conventional. Still, the discretion that conventional arbitration affords the arbitrator more than compensates the informational losses. On balance, at least in our application, conventional arbitration achieves more efficient outcomes.

Besides efficiency, we investigate how asymmetries between the disputing parties’ risk attitudes tip the scales of arbitration. Specifically, our estimates indicate that one of the parties in our empirical setting is systematically more risk-averse than the other. Our structural model enables us to assess whether such an imbalance puts

³That is, the offers in conventional arbitration are only suggestions to the arbitrator and do not affect the parties’ payoffs other than through the arbitrator’s beliefs. In contrast, final-offer arbitration has a built-in cost for aggressive offers—as, holding constant the arbitrator’s beliefs, overly ambitious offers are less likely to be selected as the ruling. This feature of final-offer arbitration makes it akin to a costly signaling game.

⁴Recently, these concerns helped motivate the choice of final-offer arbitration as the default dispute resolution method between digital platforms, such as Facebook and Google, and news outlets in Australia under the country’s News Media Bargaining Code, passed into law in February 2021. In his defense of the law, Rodney Sims, the chair of the Australian Competition and Consumer Commission, cited as the primary advantage of final-offer arbitration that “it stops ambit claims” (Senate Standing Committee on Economics, Parliament of Australia, 2021).

⁵We discuss this notion of efficiency in more detail in Section 6.3 and the relevant New Jersey statutes in Section 2.1.

the more risk-averse party at a disadvantage in arbitration. In doing so, our analysis speaks to the equity of arbitration outcomes, connecting to an ongoing debate of whether arbitration constitutes an uneven playing field for the parties involved; see, for example, Barr (2014) and Egan et al. (2018) and the New York Times article by Silver-Greenberg and Gebeloff (2015).⁶ Interestingly, we find that expected arbitration awards can actually be more favorable to parties with higher degrees of risk aversion—especially in the context of final-offer arbitration. This result arises because, in equilibrium, risk-averse parties submit more moderate offers that the arbitrator is more likely to choose. That said, we find a clear negative relationship between a party’s degree of risk aversion and its certainty equivalent of going into arbitration. That is, risk aversion makes a party worse off *ex ante* due to the associated risk premium.

Our research draws new data from the State of New Jersey, where unions must renegotiate the officers’ contracts with their employers roughly every two to three years. If the parties cannot reach an agreement, the state law requires the case to proceed to arbitration.⁷ We exploit an empirical opportunity provided by the transition of the default arbitration method from final-offer to conventional in 1996. Our data contain the parties’ offers and the arbitrator’s ruling for every case decided through final-offer arbitration between 1978-1995 and through conventional arbitration between 1996-2000. We obtain the pre-1996 final-offer arbitration data from Ashenfelter and Dahl (2012), and, as far as we are aware, ours is the first paper in the economics literature to systematically collect and investigate the post-1996 conventional arbitration data.

To analyze these data, we develop a theoretical model of arbitration that accounts for the strategic interaction between the two disputing parties—the union and the employer—and the arbitrator. The two parties are in a dispute over the wage increase, and, as in the model originally proposed by Farber (1980), we allow them to have asymmetric risk-attitudes. Additionally, motivated by evidence from the literature and following Gibbons (1988), our model accommodates learning by the arbitrator. More precisely, both the arbitrator and the disputing parties are uncertain about

⁶Most existing analyses investigate the potential disparities arising in arbitration when one of the parties is more familiar with the process or has access to better resources. These concerns are common in consumer or employment disputes between individuals and large entities such as corporations. Here, instead, we focus on disputes between organizations with comparable experience in arbitration but that might present different risk-attitudes.

⁷Per the introduction, New Jersey is not unique in relying on arbitration to resolve disputes between local governments and their employees. This procedure is especially important in disputes involving essential workers, such as police and fire officers, who are forbidden to strike by law.

what constitutes the fair wage increase, as interpreted by New Jersey law, in a given case. After filing for arbitration, the disputing parties and the arbitrator privately receive noisy signals about the fair wage increase. Next, the parties submit their offers to the arbitrator. The arbitrator employs any information about the parties' signals conveyed by the offers to update her beliefs about the fair wage increase, and then makes a decision on the case.

We bring the model to the data, initially focusing on final-offer arbitration. Specifically, we characterize the model equilibrium and formally establish identification of the model primitives under final-offer arbitration. We recover the parties' risk attitudes from the conditional odds that the arbitrator chooses the offers of one side versus the other. Intuitively, more risk-averse parties make less aggressive offers, which the arbitrator is more likely to select in equilibrium. Identification of the prior distribution of the fair wage increase and the parties' signal distribution is based on the observed joint distribution of final offers. Building upon the constructive identification argument, we propose a multi-step estimator, which we implement employing data from 1978-1995—the period when final-offer arbitration was the default arbitration procedure in our setting.

Using the estimated model, we analyze the differences between the final-offer and conventional arbitration designs by leveraging the 1996 change in the default arbitration method in New Jersey. We combine our model estimates with observed characteristics of cases decided by conventional arbitration after 1996 to simulate hypothetical outcomes of these cases under final-offer arbitration. This approach allows us to compare the two dispute resolution methods without taking a stance on the equilibria being played in the cheap-talk game implied by conventional arbitration.

We find that the expected gap between the offers made by the union and the employer more than doubles, i.e., the parties take more exaggerated positions, under conventional arbitration compared to the final-offer scenario. This result lends support to the hypothesis that the cheap-talk nature of conventional arbitration leads the parties to make offers that are not as informative to the arbitrator as those made under final-offer arbitration. To investigate this possibility in depth, we develop a new metric for information transmission in arbitration. The key idea behind the metric is to compare the observed conventional arbitration outcomes with a series of counterfactual conventional arbitration benchmarks simulated under different degrees of information transmission, which we are able to compute given our model primitives

estimated from the final-offer arbitration sample. Our results suggest that the information conveyed by the parties to the arbitrator through final offers is more than twice as precise as that transmitted in conventional arbitration; whether the game is a cheap-talk game or not is indeed consequential. But the superior information transmission afforded by final-offer arbitration comes at the cost of its one-offer-or-the-other constraint on the arbitrator's ruling. On balance, we find that conventional arbitration does better in terms of delivering arbitration awards that are closer to the ideal or fair wage. By this criterion, in our application, it is worth sacrificing the extra information of final-offer arbitration to free up the arbitrator's choice.

In a different counterfactual exercise, we shift our attention to the matter of equity between the disputing parties involved in arbitration. Specifically, we investigate how differences in risk-attitudes between the parties affect the outcomes of dispute resolution. Our baseline estimates indicate the union is risk-averse, while we let the employer be risk-neutral.⁸ As a counterfactual, we simulate a hypothetical scenario in which both parties are risk-neutral. The comparison between the baseline and counterfactual scenarios indicates that the union's risk aversion actually raises the expected salary increase for arbitrated cases, as it makes it more likely that the arbitrator chooses the union's offer in equilibrium. Nevertheless, due to the risk premium associated with the arbitrator's decision, the certainty-equivalent of going into arbitration is lower for the risk-averse union. That is, risk aversion worsens the prospects of arbitration.

In comparing conventional versus final-offer arbitration, our work pertains to the general question of how cheap-talk and costly signaling versions of a game compare empirically. Due to the nature of cheap-talk and the unobservability of private information, its empirical study has been difficult; Backus et al. (2019) remark on the paucity of empirical work on signaling games despite their theoretical importance in a wide range of domains.⁹ In particular, previous research directly comparing the information transmission in costly signaling versus cheap-talk either is purely theoretical (Austen-Smith and Banks, 2000) or employs laboratory experiments (De Haan et al.,

⁸We discuss the rationale for the risk-neutral employer in Section 2.3.

⁹Recent empirical studies on costly signaling à la Spence (1973) include Kawai et al. (2022), Sahni and Nair (2020) and Sweeting et al. (2020), whereas Backus et al. (2019) document cheap-talk signaling.

2015).¹⁰ We believe that our study is the first to undertake this type of comparison using field data.

Our paper also fits within an established literature on arbitration dating back to Stevens (1966). On the theoretical front, we contribute by characterizing the equilibrium of a final-offer arbitration model that brings together key elements from previous studies—namely, asymmetric risk-attitudes by the parties (Farber, 1980) and learning by the arbitrator (Gibbons, 1988).¹¹

Empirically, our analysis is the first to structurally estimate a model of the strategic interaction between the disputing parties and the arbitrator.¹² Our approach allows us to advance a large literature that addresses the differences between conventional and final-offer arbitration, using data from the field and lab experiments. This literature analyzes how the arbitration format affects outcomes directly observed in the data, such as the award set by the arbitrator (Bloom, 1981); the parties' willingness to make concessions and satisfaction with the dispute resolution procedure (Neale and Bazerman, 1983); and the likelihood of pre-arbitration settlement (Ashenfelter et al., 1992; Dickinson, 2004). Using our structural model, we are able to go beyond the analysis of observed outcomes to gauge the effect of arbitration design on information transmission and the efficiency of arbitration outcomes. In a similar vein, the structural approach allows us to estimate the parties' risk attitudes and disentangle their role in arbitration—a goal that has been especially elusive to empirical studies of arbitration using field data, which need to rely on proxies for the parties' risk preferences (Currie, 1989; Marburger and Scoggins, 1996).

The rest of the paper is organized as follows: Section 2 describes wage arbitration for New Jersey police and fire officers and presents the data. Section 3 contains the

¹⁰De Haan et al. (2015) consider a setup closely related to the original model by Crawford and Sobel (1982), with one privately informed sender and one receiver. Although not directly comparable to ours, their results also indicate that costly signaling allows for more informative messages.

¹¹Other theoretical studies of arbitration include Crawford (1979), Farber (1980), McCall (1990), Samuelson (1991), Farmer and Pecorino (1998), Olszewski (2011), Mylovanov and Zapechelnuk (2013), and Çelen and Özgür (2018).

¹²Ashenfelter and Bloom (1984) and Farber and Bazerman (1986) estimate a model of the arbitrator's preferences, taking the offers by the parties as exogenous. Looking at conventional and final-offer arbitration, these papers find evidence that the objective function of the arbitrators does not vary with the arbitration design. Egan et al. (2018) calibrate a model of arbitrator selection, without focusing on the strategic interaction between the parties during arbitration. Methodologically, our paper relates to a broader literature devoted to the structural analysis of bargaining and dispute resolution models. See, for example, Waldfoegel (1995), Merlo (1997), Sieg (2000), Eraslan (2008), Watanabe (2006), Merlo and Tang (2012, 2019a,b), Silveira (2017), Ambrus et al. (2018), Larsen (2020), Bagwell et al. (2020) and Larsen and Freyberger (2021).

theoretical model, and Section 4 presents our structural framework and identification results. In Section 5, we describe our estimation procedure and report the estimation results. Section 6 contains the counterfactual analyses, and Section 7 concludes. An online appendix collects proofs and supplementary analyses.

2. INSTITUTIONS AND DATA

2.1. Collective negotiations of police and fire officers in New Jersey. In 1977, the New Jersey Fire and Police Arbitration Act established a system of arbitration to avoid impasse in public sector labor negotiations. If police and fire employee unions and their municipal employers did not reach an agreement 60 days before expiry of the current labor contract, the two parties were required to file for arbitration. Until 1996, the default arbitration procedure specified by the law was final-offer arbitration. In that year, a reform instituted conventional arbitration as the new default. The reform was prompted by a perception that the final-offer arbitration design caused wages more favorable to the union,¹³ a pattern our model in Section 3 will account for.

The New Jersey Public Employment Relations Commission (PERC) oversees each arbitration case. After the disputing parties file for arbitration, PERC provides a list of seven arbitrators randomly chosen from a panel of about 60 professionals. Each party then strikes up to three names from the list, and ranks the remaining four names in order of preference. PERC then assigns to the case the arbitrator with the highest preference in the combined rankings. This selection process favors arbitrators liked by both parties. It is thus not surprising that previous studies, including Ashenfelter and Bloom (1984), Ashenfelter (1987), and Ashenfelter and Dahl (2012), find evidence that arbitrators in New Jersey are impartial and exchangeable.

The arbitration proceedings are governed by New Jersey statutory law. The law requires the arbitrator to make a decision based on a list of statutory criteria, such as the compensation currently received by the employees involved in the dispute; the continuity and stability of employment; the wages, hours and working conditions of other employees that perform comparable services in the public and private sectors; the cost of living; the financial impact of the decision on the governing unit and its residents and taxpayers; and the interests and welfare of the public.¹⁴ This last criterion, the interests and welfare of the public, is widely regarded as the most important and all-encompassing; it is the criterion to which the other criteria ultimately

¹³See Stokes (1999).

¹⁴New Jersey Statutes Title 34, Chapter 13A, Section 16.

point. As arbitrators state, the “Interest and Welfare of the Public criterion is the most significant of all statutory factors to be considered,”¹⁵ and the “interest and welfare of the public is not only a factor to be considered, it is the factor to which the most weight must be given.”¹⁶ As for what it means, this criterion is interpreted as “encompassing the need for both fiscal responsibility and the compensation package required to maintain an effective public safety department with high morale.”¹⁷

Previous empirical research suggests that the statutory concern for the impact of arbitrators’ decisions on the interest and welfare of the public is well-justified. For example, Mas (2006) provides evidence that the awards set by arbitrators in disputes between police unions and their public employers affect police performance metrics such as the crime-clearance rate and the reported crime rate. Relating to the statutory notion of an ideal award as one that promotes the interest and welfare of the public, the results in Mas suggest the existence of some “reference” or “fair” wage, below which the performance of the police deteriorates.

Several of the statutory criteria listed above refer to local conditions, of which the disputing parties are likely to have different insight than the arbitrators. For example, the union and the employer might possess specific knowledge on the fiscal state of the governing unit, the police and fire officers’ alternative job opportunities, and the local variation in the cost of living. Meanwhile, the arbitrators’ experience deciding cases in other jurisdictions affords them unique perspective regarding criteria such as the working conditions of employees performing comparable services, as well as on the proper balance of all different criteria into forming the general interest and welfare of the public. Therefore, there is ample margin for incomplete information between the arbitrator and the disputing parties about the appropriateness of different arbitration awards as per the statutes. Incomplete information of this type is a key component of the model that we develop in Section 3.

2.2. Data. We study data from the New Jersey arbitration system, which consists of two major components. The first one is the universe of final-offer arbitration cases during 1978-1995, obtained from Ashenfelter and Dahl (2012). In the remainder of the paper, we refer to this data set as ARB_F . The second component is the universe of cases decided by conventional arbitration during 1996-2000, which we collected from the PERC website. We refer to this data set as ARB_C . Both the ARB_F and

¹⁵I/M/O Passaic County and PA Local 265, IA-2022-008 (2022).

¹⁶I/M/O Seaside Park and PBA Local 182, IA-2012-022 (2012).

¹⁷I/M/O Sayreville and PBA Local 98, IA 2006-047 (2008).

TABLE 1. Summary Statistics: Final-Offer Arbitration, 1978-1995

Sample size	586	
Job type (fraction)		
Police	0.90	
Fire	0.10	
	mean	sd
Num. years covered by contract	2.1	0.7
Wage increase (% points)	7.2	1.6
Union final offer (% points)	7.8	1.8
Employer final offer (% points)	6.1	1.6
Difference in final offers (% points)	1.7	1.6
Union win rate	0.63	–

Notes: Statistics are of the ARB_F data set (explained in the text), comprising all final-offer arbitration cases during 1978-1995.

the ARB_C data sets contain, for each case, the offers made by the disputing parties, as well as the arbitrator’s decision.

The structural analysis that we present beginning in Section 4 is based on a theoretical model of final-offer arbitration. Accordingly, the ARB_F data set constitutes our estimation sample. We use the ARB_C data set only when we compare conventional and final-offer arbitration, in Section 6. In the interest of space, the current section presents only the estimation sample in more detail.

The ARB_F data consist of 586 cases after excluding observations with missing variables.¹⁸ Wages are reported as percentage increases over the previous wages, rather than in dollars terms. Table 1 provides basic summary statistics of the data. The typical observation involves a two-year contract for a municipal police department; fire contracts are fewer as many local fire departments are volunteer units. Union final offers always demand higher wages than the final offers submitted by the employer, with an average difference of 1.7 percentage points and a maximum observed difference of 12 percentage points; Appendix A Figure A1 provides a scatterplot of the final offers. At the same time, union and employer offers are positively correlated, with a correlation coefficient of 0.57. Distributions of data on offers and arbitration awards are bell-shaped and close to symmetrical, resembling normal distributions, as

¹⁸Ashenfelter and Dahl (2012) provide 620 cases with complete data on final offers. Of these, 34 cases were in municipality-years for which we could not obtain important covariates (tax base or *othermuni* information, described in Section 2.3), leading to 586 remaining cases.

seen in Appendix A Figures A2 and A3. The shapes of these distributions inform some of the parameterization choices that we make later in our structural model.

According to Ashenfelter and Dahl (2012) and their data, the disputing parties are often represented by an expert agent, such as a lawyer. This became increasingly common practice so that, by the final three years of ARB_F , both the union and the employer had an expert agent in 84% of arbitration cases. As a robustness check on the conclusions of our study, Appendix Section E provides a subsample analysis which repeats in full the counterfactual analyses of Section 6 upon restricting the estimation sample to the subset of ARB_F where both the union and the employer use expert agents. The qualitative conclusions of the subsample and full-sample analyses are the same, and the quantitative results are also similar.

2.3. Patterns in the Data and Literature. We now present patterns in our data, as well as findings from previous empirical studies of arbitration, which motivate some of the modeling assumptions of the structural analysis we present in subsequent sections. First, we investigate the relationship between realized wage increases and covariates in Table 2. Practicing arbitrators state that arbitration awards are based on the final offers submitted to arbitration and the statutory criteria mentioned above. Positions taken by the parties prior to the final offers do not factor into their award.

In light of the statutory criterion mentioning comparison to similar employees, we construct for each contract a variable *othermuni*, defined as the simple average of arbitrated salary increases of other municipalities in the same county during the most recent year available from the perspective of the case, up to a maximum of two years preceding the contract year. We also include a dummy, denoted by *otherissues*, which indicates whether the negotiations comprise any issue in addition to the workers' wages—including, for example, holiday schedules and uniform allowances.¹⁹ By New Jersey law, the scope of negotiations excludes subjects that would place substantial limits on the legislature's policy-making powers, such as pensions. To account for the financial impact on the governing unit and residents, we include the log of taxable property per capita ("tax base"), the quantile rank of median household income among New Jersey municipalities, and the credit rating assigned to municipal debt obligations by Moody's Investors' Service, as obtained from the New Jersey Data

¹⁹The ARB_F data, which we obtain from Ashenfelter and Dahl (2012), only contain the *otherissues* dummy, and do not specify at the case level what issues other than wage increases were included in the negotiations. For the ARB_C data, we observe all the negotiated issues, and find that, among the items not directly related to compensation, vacation/holiday schedules and uniform allowances are the most frequent ones.

TABLE 2. Determinants of Arbitrated Wages, 1978-1995

	(1)	(2)
Num yrs covered by contract	0.064 (0.114)	0.045 (0.101)
CPI 12 mo pct change	0.044 (0.029)	0.045 (0.025)
Othermuni	0.243 (0.053)	0.294 (0.047)
Log tax base	0.274 (0.129)	0.284 (0.094)
Income quantile	0.421 (0.306)	
Log population	-0.100 (0.066)	
Population density	0.030 (0.012)	
Fire dummy	-0.002 (0.219)	
County dummy	-0.088 (0.320)	
Otherissues	-0.077 (0.174)	
Year group fixed effects	Y	Y
Moody's rating fixed effects	Y	N
Moody's rating joint test p-value	0.50	–
Arbitrator fixed effects	Y	N
Arbitrator joint test p-value	0.94	–
Observations	579	586
R^2	0.424	0.329
Adjusted R^2	0.312	0.321

Notes: This table reports OLS results. The unit of observation is a case. In all specifications, the dependent variable is the wage increase in percentage points. Standard errors are provided in parentheses. Arbitration cases are from the ARB_F data set. See text for further details.

Book.²⁰ To account for time effects such as changes in the cost of living, we include year-group fixed effects²¹ and the 12-month percent change in the Consumer Price Index.²² Finally, we account for characteristics of the contract and bargaining units, including population as a proxy for size of the bargaining unit; a dummy indicating that the contract is for fire rather than police officers; a dummy indicating whether the employer is a county, as opposed to a municipality; and contract length in years.

Column (1) regresses arbitrated wage increases in ARB_F on these covariates. Both *othermuni* and the log tax base have a positive, statistically significant relationship with arbitrated wages. This result is consistent with intuition that arbitrators are more likely to favor higher wages if comparable employees elsewhere receive high wages and if the tax base is larger. On the other hand, other covariates such as the Moody's ratings do not have a statistically significant effect. Arbitrator fixed effects are also jointly statistically insignificant, with a p-value of 0.94. Neither do we find a significant effect for *otherissues*, indicating that the discussion of non-salary issues does not affect arbitrated wages. This result is consistent with the view by Ashenfelter and Bloom (1984) that wage increases are the focus of the disputes in this setting. Column (2) uses a more concise set of covariates, and achieves an adjusted R^2 similar to that of column (1).

Next, we investigate how choosing a higher or lower final offer affects the union's and employer's probability of winning arbitration. As the arbitrator is constrained to impose one of the two final offers in final-offer arbitration, there exists a winner by definition. We first regress union and employer final offers, respectively, on all the covariates in Table 2, column (1). We then take the respective regression residuals as a measure of how high or low each final offer is relative to the expected offer conditional on covariates. Finally, we perform probit regressions with an indicator for the employer winning as the dependent variable and these final offer residuals as the regressors. We find that a more aggressive (moderate) final offer decreases (increases) the probability of winning for both sides. Appendix A Table A1 provides detailed results. These properties shed light on the strategic considerations at play in choosing

²⁰The New Jersey Data Book only makes available the information on taxable property per capita and the municipalities' credit rating from the year 1983 onward. In our empirical analysis, we input the 1983 values of these variables for the years 1978 to 1982.

²¹There are four year-groups, 1978-1986, 1987-1990, 1991-1992, 1993-1995, formed using tests of equality of year fixed effects within groups.

²²Consumer Price Index for Urban Wage Earners and Clerical Workers in NY-NJ-PA, U.S. Bureau of Labor Statistics.

final offers; each side must trade off the gain from having a more aggressive offer accepted against the reduced probability of a more aggressive offer being accepted.

As shown in Table 1, the union wins more often than the employer. This pattern is consistent with findings by Bloom (1981) and Ashenfelter and Bloom (1984) that the union behaves conservatively in arbitration, both in an absolute sense and relative to the employer. In light of this pattern, in our structural analysis, we consider a model that allows the union to be more risk-averse than the employer. Such an asymmetric treatment of the parties' risk attitudes is not new to the literature—being adopted, for example, in papers that empirically investigate labor union preferences (Farber, 1978; Carruth and Oswald, 1985). In the public sector context, Farber and Katz (1979) explain why unions would have higher aversion to risk than their employers, stating that “wages are the primary source of income of union members, and the penalties for losing the members' primary income source are liable to be severe. On the other hand wages are not the only expense of the government unit and the taxes that finance wages account for only a small share of the expenses of the citizenry.”

Finally, the literature abounds in evidence that the parties' offers influence the arbitrator. Clearly, in final-offer arbitration, the offers directly affect the arbitrator's decision, since the arbitrator is constrained to choose one of them. But the previous literature has also provided evidence that the offers affect the arbitrator's beliefs about what the right decision should be—that is, the arbitrator learns about the case through the offers. Bazerman and Farber (1985) and Farber and Bazerman (1986) survey practicing arbitrators on hypothetical wage arbitration cases. They find that arbitrators' decisions place more weight on the parties' offers when they are of higher quality as measured by how close the two offers are. This suggests that arbitrators assess and learn from the informational content in the parties' offers. The survey responses also reveal considerable variation in arbitrator rulings given identical arbitration cases, evidencing the existence of uncertainty in arbitration outcomes. In a similar vein, Bloom (1986) conducts a survey with practicing arbitrators, asking them about hypothetical cases based on actual police wage disputes decided in New Jersey—the exact same setting of our analysis. The paper finds evidence that the parties' offers influence arbitrators' decisions in conventional arbitration. Taken together, these findings from the received literature motivate us to consider a model in which offers may convey information to the arbitrator.

3. THEORETICAL MODEL

We model two agents, a union and an employer, negotiating a wage increase, incorporating key features of the dispute resolution system described above. Henceforth, we collectively refer to the union and the employer as the *parties*. In final-offer arbitration, each party submits an offer to the arbitrator regarding the wage increase. The arbitrator then imposes one of the two offers as the wage increase. This decision is binding.

3.1. Setup. Let s represent the wage increase that would maximize the “interests and welfare of the public” as set forth in New Jersey law (refer to Section 2.1); as a short hand, we refer to this as the ideal or fair wage increase. The units of wage increase are percentage points, as our data are in percentage points. Denote by y the increase actually set by the arbitrator. The arbitrator’s utility function is $u_a(y, s) = -(y - s)^2$; that is, the arbitrator would like the expected distance between the arbitration award and the fair wage to be as small as possible.²³ For tractability, we assume a CARA specification for the union’s utility: $u_u(y) = [1 - \exp(-\rho y)] / \rho$, where the parameter ρ is common knowledge to all players. As for the employer, we assume risk-neutrality: $u_e(y) = -y$.²⁴

Neither the arbitrator nor the parties are certain about the true value of s ; as noted above, the literature finds considerable variation and uncertainty in arbitrator rulings. Instead, all players perceive s with noise. Per the description in Section 2.1, the arbitrator draws a signal of s that is separate from the parties’: the arbitrator privately receives a signal $s_a = s + \varepsilon_a$, and the parties receive a signal $s_p = s + \varepsilon_p$. Following Gibbons (1988), we let the signal s_p be common knowledge between the union and the employer. New Jersey arbitration practitioners whom we surveyed confirm that, when the parties write the final offers recorded in our data, there is no relevant information that only one side possesses, and each side is aware of what offer the other side will submit. Bloom (1981) also points out that the parties are permitted

²³Alternatively, using an absolute loss function would not affect the arbitrator’s behavior. The expected absolute loss is minimized by the median, and the quadratic loss is minimized by the mean. As we explain below, the arbitrator’s belief regarding the fair wage will have a symmetric distribution, so the median and the mean are the same here.

²⁴In addition to the reasons for a risk-neutral employer per Section 2.3, preliminary estimation allowing CARA utility for both parties yielded estimates for the employer’s risk aversion parameter that were very close to zero, as the end of Appendix C elaborates. In the text we focus on the case of a risk-neutral employer, which substantially simplifies the notation.

to adjust their offers during the course of a proceeding, so the final, observed offers reflect each party's response to the other under mutual complete information.

The incomplete information of interest in this arbitration game is between the arbitrator and the parties. The parties do not observe s_a , so they are uncertain about the arbitrator's beliefs, and neither does the arbitrator observe s_p . We make the following assumptions about the information structure:

ASSUMPTION 1. (i) The terms s , ε_a and ε_p are mutually independent; (ii) the distribution of s is normal with mean m and precision h (i.e., variance $1/h$); and (iii) the distributions of ε_a and ε_p are both normal with mean zero and precision h_ε (i.e., variance $1/h_\varepsilon$).

The normal information structure we adopt is in line with the shape of our data as discussed in Section 2.2. Though normal distributions allow negative values, our structural estimates in Section 5 indicate the proportion of the prior distribution $N(m, 1/h)$ that falls below zero is negligible in our estimated model, at about 5×10^{-5} on average.²⁵ It is also worth pointing out that, while Assumption 1 imposes that ε_a and ε_p have the same precision h_ε , we can extend our analysis to the case of asymmetric signal precision between the arbitrator and the parties. Appendix D develops such an extension. There, our estimation results indicate that the difference in precision between ε_a and ε_p is not statistically significant. In light of this finding, and accounting for expositional simplicity, we focus on the case of symmetric signal precision in the main text.

The order of play is as follows: after the parties observe s_p and the arbitrator observes s_a , the union and the employer simultaneously make final offers y_u and y_e , respectively. The arbitrator then selects either y_u or y_e as the actual wage increase.

3.2. Equilibrium. The relevant equilibrium concept is Perfect Bayesian Equilibrium. In equilibrium, the arbitrator updates her beliefs about the ideal wage increase s —based on the signal s_a , which she observes directly, and on any information about the signal s_p conveyed by the parties' final offers. Such updating by the arbitrator is consistent with the literature showing that arbitrators' opinions are influenced by final

²⁵Normality assumptions are common even when the variable in question is non-negative, both in general and especially concerning information structure. For example, the finance literature commonly models traders' information structure about stock prices as normal though negative stock prices are impossible; see, e.g., Madhavan (1992). Normality assumptions are also commonly employed in structural analyses of Bayesian learning models, as in Miller (1984), Crawford and Shum (2005) and Chan et al. (2022).

offers, as discussed in Section 2.3. She then selects the final offer that is closer to her updated expectation of s , denoted $y_a(s_a, y_u, y_e)$. That is, the arbitrator chooses the employer's offer if and only if $y_a(s_a, y_u, y_e) - y_e < y_u - y_a(s_a, y_u, y_e)$, or, equivalently,

$$y_a(s_a, y_u, y_e) < (y_u + y_e)/2 \equiv \bar{y}. \quad (1)$$

Then the union's and employer's problems in choosing final offers are, respectively,

$$\begin{aligned} & \max_{y_u} u_u(y_e) \Pr[y_a(s_a, y_u, y_e) < \bar{y}|s_p] + u_u(y_u) \{1 - \Pr[y_a(s_a, y_u, y_e) < \bar{y}|s_p]\}, \\ \text{and} \quad & \max_{y_e} u_e(y_e) \underbrace{\Pr[y_a(s_a, y_u, y_e) < \bar{y}|s_p]}_{\Pr(\text{employer wins}|s_p)} + u_e(y_u) \underbrace{\{1 - \Pr[y_a(s_a, y_u, y_e) < \bar{y}|s_p]\}}_{\Pr(\text{union wins}|s_p)}. \end{aligned}$$

The arbitrator's, union's and employer's equilibrium strategies— $y_a(s_a, y_u, y_e)$, $y_u(s_p)$ and $y_e(s_p)$, respectively—constitute a set of mutual best-responses. In particular, the final offer strategies of the union and the employer optimally balance a number of considerations: the gain from having a more aggressive offer accepted, the reduced probability of a more aggressive offer being accepted, and the opportunity to influence the arbitrator's beliefs through $y_a(\cdot, \cdot, \cdot)$. As we show below, the balance of these incentives endogenously generates divergence between the parties' positions.

By Assumption 1, Bayesian updating in this model is characterized by the normal learning model (DeGroot, 2005). Specifically, the parties' belief about the distribution of s , conditional on their signal s_p , is normal with mean

$$M_p(s_p) = \frac{hm + h_\epsilon s_p}{h + h_\epsilon}$$

and precision $h + h_\epsilon$. Also, the parties' belief about the distribution of the arbitrator's signal s_a , conditional on s_p , is normal with mean $M_p(s_p)$ and precision $H \equiv [h_\epsilon(h + h_\epsilon)] / (h + 2h_\epsilon)$. When both parties are risk-neutral, Gibbons (1988) proves the existence of a separating equilibrium in which $y_u(s_p) = M_p(s_p) + \delta$ and $y_e(s_p) = M_p(s_p) - \delta$, where δ is decreasing in the precision parameters h and h_ϵ but does not depend on the realization of s_p . That is, the union and employer strategically choose to depart from their conditional expectation of s , and the distance between their offers increases in the amount of uncertainty surrounding the case.

In Proposition 1, we show the existence of and characterize a separating Perfect Bayesian Equilibrium of our arbitration model, which allows for risk-averse or risk-loving utility and asymmetric risk attitudes between the two parties. Intuitively, final-offer arbitration has a built-in penalty for aggressive offers, as the arbitrator is less likely to choose them. This built-in penalty reins in the degree of aggressiveness

and provides for a separating equilibrium, in which the arbitrator can infer s_p from the final offers. Extending Gibbons (1988), we show that, in such an equilibrium, each party's final offer departs from $M_p(s_p)$ by a distance that depends on the precision parameters h and h_ε and the risk aversion parameter ρ , but not on the realization of s_p . This extension to asymmetric risk attitudes is not trivial because the original proof of Gibbons (1988) relies heavily on symmetry of the parties. In Proposition 2, we also show that, in this equilibrium, the distance between final offers is strictly decreasing in h and h_ε and that the more risk-averse party makes a more moderate offer, choosing a distance from $M_p(s_p)$ that is smaller than that of the opponent. All proofs of the paper are in Appendix C.

PROPOSITION 1. *Under Assumption 1, there exists a separating Perfect Bayesian Equilibrium of the arbitration game in which the final offers by the union and the employer have the form $y_u(s_p) = M_p(s_p) + \delta_u$ and $y_e(s_p) = M_p(s_p) - \delta_e$. The terms δ_u and δ_e are unique and do not depend on the signal s_p .*

To elaborate, in the equilibrium of Proposition 1, the arbitrator knows that

$$[(y_u - \delta_u) + (y_e + \delta_e)]/2 = \bar{y} + (\delta_e - \delta_u)/2 = M_p(s_p),$$

where $\bar{y} \equiv (y_u + y_e)/2$. Therefore, the arbitrator can infer s_p by applying $M_p^{-1}(\cdot)$ to both sides of the equation above, yielding the inference rule

$$s_p(\bar{y}) = \frac{(h + h_\varepsilon)[\bar{y} + (\delta_e - \delta_u)/2] - hm}{h_\varepsilon}. \quad (2)$$

This expression characterizes the arbitrator's belief about s_p , conditional on the parties' final offers, both on and off the equilibrium path. Then, given s_a and $s_p(\bar{y})$, the arbitrator updates her beliefs about s . By Assumption 1 and the normal learning model, her updated expectation of the ideal wage increase is

$$y_a(s_a, y_u, y_e) = \frac{hm + h_\varepsilon s_p(\bar{y}) + h_\varepsilon s_a}{h + 2h_\varepsilon}.$$

Then, rearranging (1), we have that the arbitrator chooses y_e if and only if

$$s_a < \frac{h_\varepsilon \bar{y} + h(\bar{y} - m) + h_\varepsilon(\bar{y} - s_p(\bar{y}))}{h_\varepsilon} = \bar{y} - \left(\frac{h + h_\varepsilon}{h_\varepsilon}\right) \frac{\delta_e - \delta_u}{2} \equiv S(\bar{y}), \quad (3)$$

where the equality comes from (2).

As previously stated, the parties' belief about the distribution of the arbitrator's signal s_a , conditional on s_p , is normal with mean $M_p(s_p)$ and precision $H \equiv [h_\varepsilon(h + h_\varepsilon)] / (h + 2h_\varepsilon)$. Denote by $\Phi(\cdot)$ and $\phi(\cdot)$ the standard normal cumulative

distribution and density functions, respectively. Then, by (3), the probability of the employer winning conditional on s_p is equal to $\Phi([S(\bar{y}) - M_p(s_p)]\sqrt{H})$. Using this expression in the union's and employer's optimization problems above, we show that the following system of first-order conditions characterizes the equilibrium values of δ_u and δ_e :

$$\frac{\sqrt{H}}{2} \frac{\phi(\eta(\delta_u - \delta_e)/2)}{1 - \Phi(\eta(\delta_u - \delta_e)/2)} = \frac{\rho}{\exp(\rho(\delta_u + \delta_e)) - 1}, \quad (4)$$

$$\text{and} \quad \frac{\sqrt{H}}{2} \frac{\phi(\eta(\delta_u - \delta_e)/2)}{\Phi(\eta(\delta_u - \delta_e)/2)} = \frac{1}{\delta_u + \delta_e}, \quad (5)$$

where $\eta \equiv \sqrt{H}(h + 2h_\varepsilon)/h_\varepsilon$. Recall that $M_p(s_p) = \bar{y} + (\delta_e - \delta_u)/2$, and δ_u, δ_e do not vary with s_p in equilibrium. By definition of $S(\bar{y})$ in (3), the probability of the employer winning is equal to

$$\Phi([S(\bar{y}) - M_p(s_p)]\sqrt{H}) = \Phi(\eta(\delta_u - \delta_e)/2) \quad (6)$$

in equilibrium. Also, taking a ratio of (4) over (5) yields

$$\frac{\Phi(\eta(\delta_u - \delta_e)/2)}{1 - \Phi(\eta(\delta_u - \delta_e)/2)} = \frac{\rho(\delta_u + \delta_e)}{\exp(\rho(\delta_u + \delta_e)) - 1}, \quad (7)$$

where the left-hand side equals the odds of the employer winning in equilibrium. We are now ready to state our next theoretical result.

PROPOSITION 2. *The equilibrium characterized in Proposition 1 is such that: (i) the distance between final offers $\delta_u + \delta_e$ is strictly decreasing in the precision parameters h and h_ε ; and (ii) the more risk-averse party chooses a final offer that is less distant from $M_p(s_p)$ —i.e., a smaller δ —and wins more often in expectation.*

The notion that the more risk-averse party wins more often in arbitration goes back to the seminal work of Farber (1980), who analyzes a simpler model in which there is no information communicated from the parties to the arbitrator. Our Proposition 2 generalizes this finding, showing that it continues to hold in an arbitration model with strategic communication.

We are aware of two existing arbitration models that characterize equilibrium offer strategies given learning by the arbitrator: Gibbons (1988) and Samuelson (1991). Samuelson (1991) proposes a model closely aligned with sealed-bid auctions where the union and employer separately receive independent private information, whereas in our model the disputing parties share the same signal that is also correlated with that of the arbitrator through s . An equilibrium implication of the Samuelson (1991)

model is that the party submitting the more aggressive or extreme offer is more likely to win, in contrast to the patterns in our data (see Section 2.3).

4. STRUCTURAL MODEL

4.1. Data Generating Process. In our structural analysis, we consider every instance of arbitration between a union and an employer as a *case*, which we index by i . We treat the precision of the signals received by the parties and the arbitrator, $h_{\varepsilon,i}$, as a random variable, which has a distribution function $G_{h_\varepsilon}(\cdot)$ and is i.i.d. across cases. We assume that the following random variables are i.i.d. across cases: the ideal wage increase, s_i ; and the noise terms $\varepsilon_{p,i}$ and $\varepsilon_{a,i}$, conditional on $h_{\varepsilon,i}$.

The model primitives are then: the union's risk aversion parameter, ρ ; the parameters of the fair wage increase distribution, m and h ; and the distribution of signal precision, $G_{h_\varepsilon}(\cdot)$. For every case, we observe the final offers by the union and the employer—respectively $y_{u,i}$ and $y_{e,i}$ —as well as y_i , the offer chosen by the arbitrator.

Our empirical analysis allows the model primitives to vary with a vector of observable case characteristics, denoted by x_i . Section 5 explains in more detail the way we account for these observable characteristics in our estimation procedure. For ease of notation, we do not explicitly condition the model primitives on x_i in our discussion of the identification strategy below. Also to facilitate the notation, we omit the index i when we refer to a specific case.

4.2. Identification. Our identification argument is constructive. A high-level intuition for it is that each h_ε is identified from the observed distance between union and employer final offers based on the monotonicity established in Proposition 2(i); the distribution of final offers conditional on between-offer difference identifies the parameters m and h ; and risk attitude ρ is identified from a conditional probability of the employer/union winning based on Proposition 2(ii).

PROPOSITION 3. *Under Assumption 1 and the equilibrium of Proposition 1, the model primitives ρ , m , h and nonparametric distribution $G_{h_\varepsilon}(\cdot)$ are identified from the joint distribution of final offers y_u and y_e and the arbitrator's decision y .*

The proof of Proposition 3 derives, among other things, the following relationship between prior precision h and the conditional variance of final offers, which we reference in the estimation section.

$$\frac{1}{H} = \left(\frac{1}{h \text{Var}[y_u|y_u - y_e]} - 1 \right) \left(\frac{1}{h} + \text{Var}[y_u|y_u - y_e] \right). \quad (8)$$

5. ESTIMATION

Our estimation procedure closely follows the identification strategy above. We accommodate observed case heterogeneity by allowing the model primitives to vary with a vector of case characteristics, denoted by x_i . This vector contains the following covariates from Table 2, column (2): the 12-month percent change in the Consumer Price Index; the log of taxable property per capita in the municipality (*log tax base*); the number of years covered by the contract; the mean arbitrated salary increase in other municipalities in the same county (*othermuni*); and year-group fixed effects. Section 2.2 provides a detailed description of each of these variables. As shown there, this set of covariates allows us to achieve explanatory power similar to that of the longer list of covariates we considered, while limiting the number of parameters to be estimated from our finite sample. Readers wishing to skip the details of implementing the estimator may proceed to Section 6 for the post-estimation analysis.

5.1. Estimation Procedure. Recall that, for every case i , we denote by $y_{u,i}$ and $y_{e,i}$ the final offers by the union and the employer, respectively. Also, define $d_{1,i} \equiv y_{u,i} - y_{e,i} = \delta_{u,i} + \delta_{e,i}$, the distance or gap between the union's and employer's final offers. Let the indicator a_i be equal to one if the arbitrator rules in favor of the employer in case i and zero otherwise.

We estimate ρ , the union's risk aversion parameter, following the argument of Proposition 3. As explained in the proof, Proposition 2(i) and (6) imply that the probability of the employer winning case i , $p_i \equiv E(a_i|h_{\epsilon,i}) = E(a_i|d_{1,i})$, is equal to $\Phi(\eta_i(\delta_{u,i} - \delta_{e,i})/2)$. Then, rearranging (7) gives

$$p_i = E(a_i|d_{1,i}) = \frac{\rho d_{1,i}}{\exp(\rho d_{1,i}) - 1 + \rho d_{1,i}}.$$

Based on this result, we propose the following estimator for ρ :

$$\hat{\rho} \equiv \arg \min_{\rho} \left[\sum_i a_i - \sum_i \frac{\rho d_{1,i}}{\exp(\rho d_{1,i}) - 1 + \rho d_{1,i}} \right]^2.$$

Next, we estimate the mean and precision of the prior distribution of the fair wage, together with the distribution of signal precision. We begin by rewriting the identifying equations in a form convenient for estimation. First, recall that, at the moment the parties formulate their final offers (that is, conditional on the parties' signal), their belief about the distribution of the arbitrator's signal has precision

$$H_i \equiv \frac{h_{\epsilon,i} [h_i + h_{\epsilon,i}]}{h_i + 2h_{\epsilon,i}}. \quad (9)$$

Plugging $p_i = \Phi(\eta_i(\delta_{u,i} - \delta_{e,i})/2)$ in (5) and rearranging yields an expression for H_i in terms of observable or known values,

$$H_i = \left[\frac{2p_i}{\phi[\Phi^{-1}(p_i)]d_{1,i}} \right]^2. \quad (10)$$

Second, rearranging (8), we obtain an expression for h_i in terms of H_i and a conditional variance of the final offers,

$$h_i = \left[\text{Var}(y_{u,i}|d_{1,i}, x_i) \left(\frac{1}{H_i} + \text{Var}(y_{u,i}|d_{1,i}, x_i) \right) \right]^{-\frac{1}{2}} \equiv \zeta_i. \quad (11)$$

Third, define $d_{2,i} \equiv (\delta_{u,i} - \delta_{e,i})/2$. Using $\eta_i \equiv \sqrt{H_i}(h_i + 2h_{\varepsilon,i})/h_{\varepsilon,i}$ and rearranging $p_i = \Phi(\eta_i(\delta_{u,i} - \delta_{e,i})/2)$ yields an expression for $d_{2,i}$,

$$d_{2,i} = \frac{h_{\varepsilon,i}\Phi^{-1}(p_i)}{\sqrt{H_i}[h_i + 2h_{\varepsilon,i}]}. \quad (12)$$

Now we set up the estimation equations. For estimation, we let the mean and precision of the fair wage depend on the covariate vector x_i according to $m_i = m(x_i; \theta_m)$ and $h_i = h(x_i; \theta_h)$, respectively, adopting the specifications

$$m(x_i; \theta_m) = x_i\theta_m \text{ and } h(x_i; \theta_h) = 1/\exp(x_i\theta_h).$$

The latter specification constrains h to be non-negative since precision is the inverse of the variance. Our task is to estimate the parameter vectors θ_m and θ_h , as well as $h_{\varepsilon,i}$, the signal precision for each case i . To estimate θ_h , let \hat{V}_i be an estimator of $\text{Var}(y_{u,i}|d_{1,i}, x_i)$,²⁶ define \hat{H}_i by substituting $\hat{p}_i \equiv \hat{\rho}d_{1,i}/[\exp(\hat{\rho}d_{1,i}) - 1 + \hat{\rho}d_{1,i}]$ for p_i in (10), and let $\hat{\zeta}_i \equiv \left[\hat{V}_i \left(1/\hat{H}_i + \hat{V}_i \right) \right]^{-\frac{1}{2}}$. Then, based on (11), we estimate θ_h as

$$\hat{\theta}_h \equiv \arg \min \sum_i \left[\hat{\zeta}_i - h(x_i; \theta_h) \right]^2.$$

We then estimate the signal precision for each arbitration case in the sample by solving for $h_{\varepsilon,i}$ in (9), using $h(x_i; \hat{\theta}_h)$ and \hat{H}_i in place of h_i and H_i . Finally, to estimate θ_m , define $\hat{d}_{2,i}$ by substituting $\hat{h}_{\varepsilon,i}$, \hat{p}_i , \hat{H}_i and $h(x_i; \hat{\theta}_h)$ for $h_{\varepsilon,i}$, p_i , H_i and h_i in (12), respectively. Then, in light of $(y_{u,i} + y_{e,i})/2 - d_{2,i} = M_p(s_{p,i})$ and $E[M_p(s_{p,i})|x_i] =$

²⁶We obtain \hat{V}_i by, first, using single index kernel regressions of the union's final-offer on $d_{1,i}$ and x_i to compute estimates of $E[y_{u,i}|d_{1,i}, x_i]$ and $E[y_{u,i}^2|d_{1,i}, x_i]$, and then applying the standard expression of the variance of a random variable in terms of the mean of its square and the square of its mean.

$m(x_i; \theta_m)$ (see Proposition 1 and the proof of Proposition 3), we estimate θ_m as

$$\hat{\theta}_m \equiv \arg \min_{\theta_m} \sum_i \left[\frac{y_{u,i} + y_{e,i}}{2} - \hat{d}_{2,i} - m(x_i; \theta_m) \right]^2.$$

5.2. Estimation Results. We now discuss our estimates of ρ , θ_m , θ_h , and $h_{\varepsilon,i}$. Our estimate of the risk aversion parameter is $\hat{\rho} = 0.60$. By definition, the CARA risk aversion parameter has units of 1/(unit of the argument). Since the argument of the utility function in our setting has units of percentage points, a comparison to measures of CARA risk aversion in other settings requires a conversion. For example, if one percentage point of wage increase represents about \$500, our CARA parameter converts to about $0.60/500 = 0.0012$ in units of 1/\$. This amount is in the range of CARA estimates from various studies summarized by Babcock et al. (1993). In the subsample analysis of Appendix E, we re-estimate the model using only observations in which both parties employed expert agents. In that analysis, we also estimate the union to be risk-averse, albeit with a smaller parameter, $\hat{\rho} = 0.32$. We find that the qualitative conclusions of Section 6 do not differ between the subsample and full-sample analyses, and the quantitative conclusions are also similar.

Next, Table 3 reports the estimates of θ_m and θ_h . For $m(x_i; \theta_m)$, we extend x_i by including the square of the number of years covered by the contract to allow for a nonlinear effect. Inflation and *othermuni* both have significant positive marginal effects on the mean m of the fair wage increase, while the effect of contract length on m is statistically insignificant. This is consistent with the patterns presented in Table 2 of Section 2.3. While the components of $\hat{\theta}_h$ are not statistically significant at conventional levels, longer contracts are associated with smaller variance, suggesting that the range of wage increases considered appropriate is narrower when the contract has longer-term influence on wages.

The median of $m(x_i; \hat{\theta}_m)$, the prior mean of the fair wage increase, is 7.5 percentage points in the ARB_F data set, while the 1st and 99th percentiles are 4.4 and 9.4 percentage points, respectively. The median of $\sqrt{1/h(x_i, \hat{\theta}_h)}$, the prior standard deviation of the fair wage increase, is 1.7 percentage points, while the 1st and 99th percentiles are 0.6 and 2.8 percentage points, respectively. Figure 1 plots the kernel density of $\sqrt{1/\hat{h}_{\varepsilon,i}}$, the estimated standard deviation of the noise term ε in the players' signals of the fair wage increase. The median of $\sqrt{1/\hat{h}_{\varepsilon,i}}$ is 0.4 percentage points, so the variance of the signal noise is typically a fraction of the prior variance of the fair wage itself.

TABLE 3. Parameter Estimates in $m(x_i; \theta_m)$ and $h(x_i; \theta_h)$


x_i	$\hat{\theta}_m$	$\hat{\theta}_h$
CPI 12mo pct change	0.11 (0.05)	0.08 (0.15)
Log tax base	0.04 (0.12)	0.01 (0.26)
Num years covered by contract	-1.05 (1.13)	-0.42 (0.33)
Squared num years covered by contract	0.17 (0.23)	-- --
Othermuni	0.34 (0.09)	0.03 (0.19)
Year group dummy, 1987–1990	-0.22 (0.19)	-0.27 (1.48)
Year group dummy, 1991–1992	-0.76 (0.22)	0.09 (1.20)
Year group dummy, 1993–1995	-1.62 (0.32)	-1.44 (1.59)
Constant	5.48 (1.85)	1.14 (2.77)

Notes: Table reports estimates of the parameters, θ_m and θ_h , of the prior mean m and precision h of the fair wage distribution. Units are percentage points of initial wages. The parentheses report standard errors computed from $B = 200$ bootstrap samples drawn from ARB_F .

To assess model fit, we perform Monte Carlo simulations with our estimated model to simulate 1000 cases for each set of covariates x_i observed in the relevant data. Figure A3 in Appendix A plots the observed versus model-simulated outcome distributions. The model achieves a close fit to the observed distribution of final offers for both the union and the employer. The model-simulated likelihood that the employer wins arbitration matches the observed employer win rate, at 0.37.

5.3. Selection into Arbitration. The ARB_F data consists of arbitrated cases, and, accordingly, our structural model abstracts away from any pre-arbitration negotiations and from the parties' decision of whether to settle their dispute. One interesting question is whether our results might be affected by the endogenous selection of cases into arbitration. In particular, if the disputing parties have access to their signal s_p at the pre-arbitration negotiations, then the realization of that signal could affect the odds that bargaining breaks down and the case reaches arbitration. This could

FIGURE 1. Density of $\sqrt{1/\hat{h}_{\varepsilon,i}}$, the Standard Deviation of Signal Noise



Notes: Figure displays kernel density of $\sqrt{1/\hat{h}_{\varepsilon,i}}$ based on Gaussian kernels and bandwidth given by Silverman’s rule of thumb. The plot is truncated at the 95th percentile.

lead to differences in the signal distribution between the cases in our sample and the general population of cases.

Appendix F explores this question by modeling pre-arbitration negotiation. The union and the employer engage in a bargaining game with two inputs: (i) the certainty equivalents of their respective arbitration payoffs conditional on the realization of their signal s_p ; and (ii) any other costs or benefits of settling versus going to arbitration, which we refer to as “backing-down costs” and assume to be independent of (i). Given this framework, we provide analyses of two distinct bargaining models that differ in whether the backing-down costs are common knowledge and also differ in bargaining protocol. For both bargaining models, we establish that what matters for the probability of settlement failure (i.e., arbitration) is not any party’s individual value of arbitration per se, but the *difference* between the two parties’ certainty equivalents of arbitration. Then, we follow this theoretical analysis with empirical work assessing how the parties’ signal s_p affects that gap in the certainty equivalents. With our estimated arbitration model, we find there is a gap between the union’s and employer’s certainty equivalents of arbitration; but, conditional on case covariates, such a gap is essentially constant in the realization of the signal s_p . It follows that the conditional probability of going to arbitration is invariant to changes in the signal s_p . This in turn implies there would not be systematic bias in the signal structure we estimate using arbitrated cases.

6. COUNTERFACTUAL ANALYSES

Having estimated our model, we now turn to addressing questions about the properties of arbitration in practice. Sections 6.1-6.3 compare the two forms of arbitration—final-offer and conventional—in terms of the offers they elicit from the disputing parties; the arbitrated outcomes; their conduciveness to information revelation; and their efficiency, as measured by the distance between arbitrated awards and the ideal or fair wage. Lastly, Section 6.4 investigates the potential inequities generated by asymmetric risk attitudes in arbitration.

6.1. Offers and awards in CA versus FOA. In this section, we compare two commonly employed forms of arbitration, final-offer (FOA) and conventional (CA), in terms of the offers they induce from the disputing parties and the resulting arbitration awards. We complement observational comparisons of FOA and CA jurisdictions and cases, such as Feuille (1975), Bloom (1981) and Ashenfelter and Bloom (1984), by leveraging our structural model to compare how the same case would fare under FOA versus CA. Specifically, we compare outcomes observed under New Jersey's implementation of CA after 1996 to counterfactual model simulations of FOA for the same arbitration cases.

Whether the offers in CA differ from those in FOA is an empirical question. Unlike FOA, where the parties' offers directly affect payoffs because one of them must be chosen as the arbitration award, CA does not impose such a constraint. As a result, the parties' offers in CA may matter only indirectly through the information they convey to the arbitrator. In other words, the offers in CA are cheap-talk. Gibbons (1988) shows that if the arbitrator in CA enforces a large transfer from the party who seems to have made the less reasonable offer to the party who seems to have made the more reasonable offer—effectively mimicking the incentives toward reasonable offers created in FOA—then there is a separating equilibrium of CA that generates the same offers as FOA. However, like all cheap-talk games, that CA game has a continuum of payoff-equivalent separating equilibria that differ only by a translation, in which the distance between parties' offers are different from those in FOA. Moreover, we have no reason to believe that arbitrators enforce such transfers in practice. The effect of FOA versus CA on the distribution of arbitrated wages is also an empirical question. On the one hand, the pendulum nature of FOA, which forces the arbitrator to choose one party's offer or the other, may increase the variance of awards by eliminating awards in the middle. On the other hand, this restriction of FOA may also serve to

eliminate the tails of potential awards and thus decrease variance, especially if the two parties' offers are closer together in FOA than in CA.

Since cheap-talk games raise the possibility that the equilibrium in play may not be separating, we do not posit any specific equilibrium for CA in our analysis. Instead, we simply report the observed outcomes of conventional arbitration in the ARB_C data set, defined in Section 2.2. We do make the following two assumptions that provide minimal structure for a meaningful comparison. The first is that in CA the arbitrator sets the award equal to y_a , her updated expectation of the fair wage after observing the offers, as defined in Section 3.2. Recall that, in FOA, the arbitrator chooses the offer that is closest to y_a as the award because the rules constrain her to choose one of the parties' offers. CA does not impose such constraints and gives the arbitrator freedom to impose y_a directly.²⁷ This is useful because it implies that we observe y_a directly from the arbitrator's awards in the ARB_C data. The second assumption is that $E[y_a] = m$ in CA, as it is in FOA. A benefit of this assumption for our analysis is that it allows us to estimate a new mean for the prior distribution of fair wage increases specifically for the post-1996 era, using the observed y_a ; with such added flexibility, we prevent our comparison between FOA and CA from being confounded by changes over time in the prior mean of fair wage increases. We can prove that our second assumption is true both in the case of a separating equilibrium and in the opposite case, when the arbitrator cannot infer any information from the parties' offers.²⁸

Implementation. As defined in Section 5, let x_i refer to covariates that describe case i . We take the following steps to minimize confounding factors when simulating FOA outcomes corresponding to each observed CA case i . First, to account for potential changes in the prior mean of fair wage increases after 1996, we specify $m_i = m(x_i; \theta'_m)$ in simulations, where θ'_m is newly estimated from post-96 data which consist of CA cases only. Specifically, since we observe arbitration awards y_a in CA, and $E[y_a] = m$, we estimate θ'_m as $\hat{\theta}'_m \equiv \arg \min_{\theta'_m} \sum_i [y_{a,i} - m(x_i; \theta'_m)]^2$. Second, recall that one of the covariates in x_i is a year-group dummy that accounts for changes across time in the estimation sample that are not already reflected in other covariates. When defining

²⁷Indeed, that the arbitrator imposes her notion of the fair wage as the award is the standard view of arbitrator behavior in conventional arbitration; see, for example, Ashenfelter et al. (1992).

²⁸In a separating equilibrium where the arbitrator infers s_p from the parties' offers, $y_a = (hm + h_\epsilon s_p + h_\epsilon s_a)/(h + 2h_\epsilon)$ by the normal learning model. In an equilibrium where the arbitrator infers nothing about s_p , $y_a = (hm + h_\epsilon s_a)/(h + h_\epsilon)$. By the definitions of s_p and s_a in Section 3, it follows immediately that $E[y_a] = m$ in both cases.

TABLE 4. Conventional Versus Final-Offer Arbitration, 1996-2000

	(1) Conventional (observed)	(2) Final-offer (simulated)	(1)-(2) 95% C.I.
(a) Mean difference between offers	2.48 (0.13)	0.92 (0.07)	[1.42, 1.68]
(b) Mean arb. wage – offer midpoint	-0.26 (0.07)	0.08 (0.02)	[-0.38,-0.31]
(c) Probability of union win	n/a	0.57 (0.01)	[-0.10,-0.05]
(d) Mean arbitrated wage increase	3.70 (0.06)	3.75 (0.02)	[-0.10,-0.02]

Notes: Column 1 shows average outcomes of the 119 observations in ARB_C . The parentheses in Column 1 report the standard errors of these sample means and proportion from ARB_C . Column 2 Monte Carlo simulates the final-offer arbitration model 1,000 times conditional on each set of covariates in ARB_C ; thus, it presents average outcomes across a total of 119,000 simulated cases. The parentheses in Column 2 report standard errors for these outcomes computed from 200 bootstrap samples of ARB_F . Column 3 reports the 95% confidence interval of the difference between the two columns (Column 1 - Column 2), using its empirical distribution from the bootstrap samples. (In row (c), column 3 shows the 95% confidence interval of 0.5-(2).) Offers and wage increases are in units of percentage points.

that dummy variable for CA cases, we group the CA years (1996-2000) only with the last year-group in the estimation sample (1993-1995), so the $h_i = h(x_i; \hat{\theta}_h)$ and $\hat{G}_{h_\epsilon}(\cdot)$ used in simulation reflect conditions of the mid-late 1990s as opposed to earlier years. This is an extra precaution; in practice, Table 3 shows that, conditional on covariates, the difference in h_i across year groups is not statistically significant.

Given these model parameters, we perform counterfactual simulations of the FOA model, 1000 times for each set of covariate values x_i observed in the ARB_C sample. The simulation process involves taking random draws of $h_{\epsilon,i}$, s_i , $\epsilon_{p,i}$, and $\epsilon_{a,i}$ conditional on the covariates and simulating the parties' final offers and arbitrator's decision.

Results. Table 4 highlights the differences we find between CA and FOA. The second column of Table 4 presents the results of the FOA simulations, while the first column presents observed CA statistics for comparison. The third column shows the 95% bootstrap confidence interval of the difference between each observed CA statistic and the simulated FOA analog; this is obtained by drawing $B = 200$ bootstrap samples from ARB_F and repeating the estimation procedure and counterfactual simulations for each bootstrap sample.

First, Table 4, row (a) shows that the gap between parties' offers is significantly narrower in FOA than in CA; in other words, the parties take more reasonable positions in FOA. Since the arbitrator is constrained to choose one of the two offers in FOA, there is pressure for the parties to submit reasonable offers in order to be the one chosen. CA offers, meanwhile, diverge more, notwithstanding the theoretical possibilities discussed above. Second, in row (b) of Table 4, we find that on average the arbitrated wage would be higher than the midpoint of offers in FOA while it is lower in CA. This difference is statistically significant and is driven by the winning offer being imposed without compromise in FOA while the union wins more than half of the time (per row (c)); the interaction of arbitration format with the union's risk aversion has consequences here. Finally, Table 4, row (d) says that, given identical cases, FOA would have yielded slightly higher arbitrated wage increases on average. The difference is statistically significant though small in magnitude. As a means of supplementing and corroborating the findings in Table 4, we present in Appendix B a descriptive regression analysis comparing cases decided by FOA during 1993-1995 and cases resolved by CA during 1996-2000. The regression results are largely consistent with the findings from the counterfactual simulation in the present section, despite methodological differences and the distinct samples used in the two analyses.²⁹ The likeness of the two sets of results provides additional reassurance regarding the robustness of Table 4, the credibility of our structural analysis in general and of the counterfactual exercises in the next sections that are motivated by these comparisons.

One caveat in interpreting Table 4 is that this comparison of FOA versus CA holds fixed the set of cases. In other words, the comparison asks how arbitration design affects offers and awards *conditional* on the same set of cases being arbitrated. We may subsequently discuss what this comparison implies about the relative attractiveness of the two designs to the parties and any consequent differences in the propensity to resort to arbitration. In particular, Stevens (1966) argues that arbitration would be less frequent in FOA because FOA generates more uncertainty for the parties, lowering a risk-averse party's certainty equivalent of arbitration. Indeed, the standard deviation of the arbitrated wage increase is 0.70 percentage points in our FOA simulations versus 0.62 in CA. However, the mean arbitrated wage increase is also slightly higher in FOA due to the union winning more than half of the time. Given the union's estimated risk aversion parameter $\rho = 0.60$, the difference in the certainty

²⁹The most different result is for Table 4, row (d); both analyses indicate higher arbitrated wage increases for FOA, but Table 4 is more conservative in its estimated effect of FOA compared to a before-and-after regression. Appendix B provides further discussion.

equivalent of FOA versus CA is less than 0.1 percentage point in the end,³⁰ conditional on the same set of cases as in Table 4. Thus, in this application, we do not find much support for Stevens’ prediction.³¹ Statistics on the number of arbitration awards before and after 1996 bear this out. Stokes (1999) reports that “the number of awards rendered under the act has not changed very much since the amendments were passed.” Our annual count of arbitration awards surrounding the policy change, displayed in Figure A4, corroborates Stokes’ report. So neither our model-based computations nor the case-count statistics bear out a substantial difference in arbitration frequency between FOA and CA in NJ. Nonetheless, there could be other differences in the set of cases that would be arbitrated.

One of the most notable results in this section is that the disputing parties’ offers are more distant in CA than in FOA, meaning that the parties take more exaggerated positions. While this does not necessarily imply that offers in CA are less informative to the arbitrator as signals of the fair wage, it is nonetheless suggestive in that regard. We investigate this possibility in the next section.

6.2. Information transmission in CA versus FOA. As explained above, a key difference between the final-offer (FOA) design and the conventional arbitration (CA) design is that the latter is a cheap-talk game, in which it may be difficult for the arbitrator to infer information about the parties’ private signal from their offers. Our estimated model of FOA combined with observed data on CA grants us a unique opportunity to assess the degree of information transmission in CA relative to FOA in practice.

For a tractable analysis, we first develop a concise representation of the degree of information transmission. Specifically, we represent the degree of information transmission by a scalar $\alpha \in [0, 1]$, where a higher value of α indicates better transmission; $\alpha = 1$ represents full communication or a separating equilibrium, $\alpha = 0$ represents no communication, and $\alpha \in (0, 1)$ represents the spectrum of imperfect information

³⁰Given that we are agnostic about the specific equilibrium in CA, we numerically approximate the union’s certainty equivalent of CA by two separate methods: 1) fitting the observed distribution of CA awards with a normal distribution and applying the analytical approximation based on normal distributions, $CE(y) = E(y) - 0.5\rho\text{Var}(y)$; and 2) exploiting the degree of information transmission we estimate in Section 6.2. Both methods yield a CA-FOA difference of less than 0.1 percentage point.

³¹In Appendix F, we formalize the argument that, under various models of pre-arbitration negotiation between the union and the employer, the probability that the parties enter arbitration remains constant if the parties’ certainty equivalents of going to arbitration do not change.

transmission in between. To aid intuition, the next paragraph provides one possible rationale for such a representation.

Recall that we denote by s_p the signal about the fair wage increase received by the parties at the beginning of the arbitration game. Suppose the arbitrator is unable to infer s_p perfectly from the arbitration process and can only infer a noisy measure of it, $s_p^* \equiv s_p + \epsilon_n$, where ϵ_n is an exogenous, mean-zero error that is normally distributed with precision h_n . Then, $s_p^* = s + \epsilon_p + \epsilon_n = s + \epsilon_p^*$, where $\epsilon_p^* \equiv \epsilon_p + \epsilon_n$ is normally distributed with mean zero and precision

$$h_p^* \equiv h_\epsilon \frac{h_n}{h_\epsilon + h_n}$$

by the Bienaymé formula for variance. The effective precision h_p^* of the signal the arbitrator infers, s_p^* , equals the original precision h_ϵ multiplied by a fraction $h_n/(h_\epsilon + h_n)$. This fraction goes to 1 as $h_n \rightarrow \infty$, the scenario in which the arbitration process perfectly reveals s_p , and goes to 0 as $h_n \rightarrow 0$, the scenario in which the arbitration process reveals nothing about s_p . Thus, on an aggregate level, we may reasonably represent the degree of information transmission by a scalar $\alpha \in [0, 1]$ so that $h_p^* = \alpha h_\epsilon$, where a higher value of α indicates better transmission.

Now consider the implications for the arbitrator's preferred award y_a as α increases. Intuitively, the more precisely the arbitrator is able to learn about s_p , the more weight she will give to it in forming her preferred award y_a . Therefore, we would expect more of the variance of y_a to be explained by s_p^* when α is larger.³² Indeed, our simulation results, to be discussed below, verify this numerically.

Thus, as an intuitive measure of information transmission, we consider the R^2 of regressing the arbitrator's preferred award, y_a , on the signal she infers from the parties' offers, s_p^* . That is, we can assess the degree of information transmission in the observed conventional arbitration (CA) data by comparing the R^2 of such a regression to that in simulated data. Specifically, we simulate y_a and s_p^* data given each *hypothetical* value of α over a grid in $[0, 1]$; then we look for the value of α , or degree of information transmission, that generates the R^2 most consistent with the R^2 obtained from the actual data. Note that we do not need to know the parties' equilibrium offer strategies in CA to be able to simulate the regressand y_a ; as before, we remain agnostic in that regard. Regardless of how she does it, if the arbitrator

³²Let \tilde{y}_a be the linear projection of y_a on s_p^* . Given the normal learning model, we can prove analytically that $\text{var}(\tilde{y}_a)/\text{var}(y_a)$ is strictly increasing in the degree of information transmission, α .

ultimately infers s_p^* as defined above, and this has precision $h_p^* = \alpha h_\epsilon$, then it follows from the normal learning model that $y_a = (hm + h_p^* s_p^* + h_\epsilon s_a)/(h + h_p^* + h_\epsilon)$.

Implementation. Given this conceptual framework, we implement our assessment as follows. First, we use our estimated model primitives to establish a mapping between all possible values of α and the R^2 described above, so that upon observing an R^2 we can interpret the implied α . We start by simulating, given each hypothetical value of α on a grid in $[0, 1]$, 1000 Monte Carlo samples of $s_p^* \equiv s + \epsilon_p^*$ and $y_a = (hm + h_p^* s_p^* + h_\epsilon s_a)/(h + h_p^* + h_\epsilon)$ per each set of covariates x_i observed in ARB_C . Note that it would not have been possible to simulate these without our estimated model of the information structure; we use the same $m_i = m(x_i; \hat{\theta}'_m)$, $h_i = h(x_i; \hat{\theta}'_h)$ and $\hat{G}_{h_\epsilon}(\cdot)$ used in the Section 6.1 simulations and described in detail there. The s and s_a are Monte Carlo simulated given these model parameters. As explained above, ϵ_p^* is normally distributed with mean zero and precision $h_p^* = \alpha h_\epsilon$, where h_ϵ is drawn from the distribution $\hat{G}_{h_\epsilon}(\cdot)$. Then using the entire Monte Carlo sample associated with each α value, we run the OLS regression

$$y_{a,i} = \beta_0 + \beta_1 m_i + \beta_2 s_{p,i}^* + \nu_i \quad (13)$$

and obtain the resulting $R^2(\alpha)$. The regressor $m_i = m(x_i; \hat{\theta}'_m)$ is simply a control for the heterogeneity of covariates across cases.

Second, we run an analogous regression using the observed CA data to obtain the relevant R^2 thereof. We will subsequently interpret the α implied by this R^2 using the α -to- R^2 mapping established through the simulations in the preceding paragraph. Here, we observe the regressand y_a directly in the data, since y_a corresponds to the observed arbitration award in CA. We also observe the offers of the two parties, but, in CA, we do not know the functional form by which they convey s_p^* . What we do know is that s_p^* is by definition something the arbitrator infers from the offers, so it is some (unknown) function of the offers. Therefore, we substitute the regressor s_p^* in regression (13) with bivariate thin plate regression splines of the observed offers of the parties. The smoothing parameter is optimized by generalized cross validation.³³ We also substitute the regressor m in regression (13) with the covariates listed in Table 2 that are available for ARB_C as well as year and credit rating fixed effects. If the observed CA data, despite generous inclusion of regressors, achieves a lower R^2 than that simulated for full information transmission, that finding would be more


³³We use the `mgcv` package in R.

indicative of weak information transmission in CA than it would be if we had not been so generous. This one regression using observed data leads to an R^2 of 0.50.

Note that, by construction, OLS regression yields the smallest possible sum of squared residuals, or highest R^2 , among all possible inference rules by the CA arbitrator that are a function of the observed offers, which we leave as unspecified in the model and remain agnostic about. So this R^2 represents the best-case scenario in terms of information transmission in conventional arbitration. Specifically, it corresponds to an inference rule by the arbitrator that leads to the highest possible R^2 in the regression of observed CA awards.

Results. Figure 2 plots the R^2 from the simulated data as a function of α using a solid curve. The monotonic increase of the curve as a function of α numerically confirms our intuition that more of the variance of y_a is explained by s_p^* when α is larger. The R^2 for the observed conventional arbitration (CA) data, 0.50, is marked by a dotted line. This observed R^2 is closest to that of the simulation in which $\alpha = 0.33$. To directly interpret this α , recall that the signal the arbitrator infers from the parties' offers is a noisy signal of the ideal or fair wage. If $\alpha = 0.33$, this means the variance of the noise is $1/0.33 \approx 3$ times larger in CA than in FOA. A 95% confidence interval for α , which is constructed from the empirical distribution of bootstrap estimates by resampling from ARB_F , is $[0.01, 0.55]$. Thus, our simulations suggest that conventional arbitration does communicate some private information from the parties to the arbitrator in a statistically significant way, though the lower bound of the confidence interval is close to zero. However, the transmitted information is also significantly less precise than that in final-offer arbitration, which is represented by the benchmark of $\alpha = 1$. In contexts where communication of private information from the disputing parties to the arbitrator is particularly important, final-offer arbitration may indeed be preferable to conventional arbitration.

The metric α is consistently far from 1 in a number of robustness analyses we conduct. First, the finding is robust to using a different type of spline. Tensor product splines are an alternative among splines that accommodate bivariate functions; using these instead of thin plate regression splines yields $\alpha = 0.30$. Second, the finding is robust to a subsample analysis in which we restrict the estimation sample to only those cases where both the union and the employer were represented by expert agents; subsequently redoing the entire analysis yields $\alpha = 0.27$ (see Appendix E). Third, the estimated α is even smaller in a symmetric utility specification: when we estimate the



Notes: Figure displays simulated R^2 values of regression (13) as a function of α , the degree of information transmission. At each value of α , we Monte Carlo simulate 1000 cases per each set of covariates observed in ARB_C and run the regression. For comparison, the dotted, horizontal line marks the R^2 of a regression analogous to (13) run using the observed data from ARB_C .

arbitration model specifying both parties as risk neutral and redo the entire analysis, we obtain $\alpha = 0.17$.

6.3. Efficiency of awards in CA versus FOA. As a final criterion of comparison, we consider the ability of each arbitration design to yield awards that are close to the fair wage increase s . Recall that s is defined as the wage that would maximize the “interests and welfare of the public” as set forth in New Jersey law. We call this criterion “efficiency” and measure it by the arbitrator’s objective function $u_a(y, s) = -(y - s)^2$. Our structural model primitives, including the distribution of fair wage increase s , allow us to assess efficiency through this criterion despite s being unobserved.

As we saw in the previous section, FOA transmits more precise information from the parties to the arbitrator than CA. However, this comes at the cost of the one-offer-or-the-other constraint on the arbitrator in FOA, which may constrain the award away from the fair wage s even while the arbitrator is better informed of what this fair wage is. Determining which arbitration design is more efficient on balance is an empirical question, not a theoretical one; numerical simulations of our theoretical model at different combinations of parameter values show that the efficiency ordering of CA versus FOA depends on the model primitives. As we do not have a closed-form

TABLE 5. Efficiency of Awards in CA and FOA

	(1) Conventional ($\alpha = 0.33$)	(2) Final-offer	(1)-(2) 95% C.I.
$E[-(y - s)^2]$	-0.06 (0.04)	-0.21 (0.04)	[0.01,0.24]
$E[- y - s]$	-0.19 (0.04)	-0.35 (0.04)	[0.04,0.28]

Notes: The table displays the mean of the efficiency measure across 1000 Monte Carlo simulations conditional on each set of covariates in the ARB_C data set; thus, it presents average outcomes across a total of 119,000 simulated cases. Standard errors in the parentheses are computed using $B = 200$ replications of bootstrap samples. Column 3 report 95% confidence intervals of the difference (Column 1 - Column 2), using the empirical distribution from bootstrap samples.

solution for final offers in FOA, it is difficult to generate an analytical characterization of when CA dominates FOA and vice versa. Nonetheless, we observe some patterns in the numerical simulations. If the prior is already very precise, more learning is not that useful to the arbitrator, so CA tends to be more efficient than FOA. At the other extreme, if there is a lot of general uncertainty (i.e., both the prior and the signals are very imprecise), this leads to a large union-employer offer gap in FOA. Then allowing compromise awards as in CA tends to be more efficient than forcing a choice between extremes as in FOA. In some intermediate cases outside of the above scenarios, FOA can dominate CA by providing sufficient added information about the location of the fair wage to overcome the disadvantage of the award constraint.

To assess which arbitration design is more efficient in New Jersey, we use the estimated structural model to numerically compare the mean of $-(y-s)^2$ across Monte Carlo simulations of FOA and CA. Specifically, for FOA we use the FOA sample simulated in Section 6.1, and for CA we use the CA sample simulated conditional on $\hat{\alpha} = 0.33$ in Section 6.2; i.e., we simulate CA given the estimated degree of information transmission. Both of these samples are conditioned on the set of covariates observed in ARB_C and are of equal sample size.

Table 5 displays the measure of efficiency thus simulated in CA versus FOA. We find that CA is more efficient; the average distance of the award from the fair wage, in terms of squared percentage points, is 0.06 in CA compared to 0.21 in FOA. Using an alternative metric, such as the absolute value of the difference between the award and the fair wage, leads to the same qualitative result. Two-sided t-tests using bootstrap

standard errors reject the null of equal efficiency loss under conventional and final-offer arbitration at the 5% significance level. Finally, we find that CA is more efficient than FOA for all α values within the 95% bootstrap confidence interval for α .

These results imply that the gain in efficiency from the arbitrator not being constrained in CA outweighs the loss in efficiency from inferior information transmission. So, as far as efficiency is concerned, it is worth sacrificing information here to free up the arbitrator’s choice. If we interpret s as the outcome that would maximize the “interests and welfare of the public” criterion specified in New Jersey law and $-(y - s)^2$ as measuring closeness to that outcome, then, by this measure, CA would be the better choice over FOA in New Jersey’s public sector labor disputes.

6.4. Asymmetric risk attitudes and (in)equity in arbitration. According to estimates from Section 5 and consistent with evidence in Section 2.3, New Jersey police and fire unions are risk-averse in the period that we analyze. Risk aversion is likely to be present in labor negotiations of other states and industries as well as in contexts other than labor, such as the arbitration of disputes between consumers and businesses. As such, we investigate how risk aversion interacts with the dispute resolution mechanism to affect arbitration outcomes. We focus on the FOA design, for which Proposition 2 established the relationship between risk attitudes and equilibrium final offers.

To study this question, we counterfactually simulate a scenario in which both the union and the employer are risk-neutral. Specifically, we perform Monte Carlo simulations of the FOA arbitration model, 1000 times for each set of covariate values x_i observed in the ARB_F data set. This results in a total of 586,000 simulated cases.

Table 6 compares simulated outcomes when the union is risk-averse, with $\rho = 0.60$ as estimated in our data, to the simulated counterfactual outcomes when the union is risk neutral. To gain a fuller view of the effects of risk aversion, the table also displays counterfactual outcomes when the union is more risk-averse than estimated in our data, with $\rho = 1.5$, but still within the range of CARA estimates reported by Babcock et al. (1993). The employer remains risk-neutral throughout. Table 6, row (a) shows that, when the union is risk-averse, it chooses a more moderate final offer than in the risk-neutral scenario, asking for a smaller wage increase. The employer is also less aggressive in response, but its offer does not change as much as the union’s. As a result, the risk-averse union wins more than half of the time, whereas both parties win with equal frequency when the union is risk-neutral. In fact, Table 6, row (d) shows that the risk-averse union obtains a slightly larger arbitrated wage increase,

TABLE 6. Risk-Averse Versus Risk-Neutral Union in FOA, 1978-1995

	risk neutral	$\rho = 0.60$	$\rho = 1.5$
(a) Mean union offer	8.73 (0.26)	8.05 (0.15)	7.70 (0.16)
(b) Mean employer offer	6.00 (0.13)	6.36 (0.16)	6.41 (0.13)
(c) Probability of union win	0.50 (0.00)	0.63 (0.02)	0.72 (0.01)
(d) Mean arbitrated wage increase	7.37 (0.16)	7.57 (0.17)	7.46 (0.16)
(e) Union's certainty equivalent	7.37 (0.16)	6.56 (0.16)	5.58 (0.23)

Notes: The FOA model is Monte Carlo simulated 1000 times conditional on each set of covariates in the ARB_F data sets; thus, the table presents average outcome across a total of 586,000 simulated cases. Units are percentage points, excluding probabilities. Employer is risk neutral throughout. Standard errors in the parentheses are computed using $B = 200$ replications of bootstrap samples.

on average, than it would in the risk-neutral scenario. This difference is statistically significant, as the 95% confidence intervals for the difference between the two risk-averse cases and the risk-neutral case—which are constructed using the empirical distribution of bootstrap estimates—are $[0.16, 0.26]$ and $[0.07, 0.16]$ respectively.

It might, at first glance, seem counter-intuitive that a risk-averse union would tend to secure higher arbitrated wage increases than a risk-neutral one. This result is due to two aspects of the parties' equilibrium behavior mentioned above. The first one is the risk-averse union's increased odds of winning arbitration. The second one relates to the relatively conservative equilibrium offer made by the employer when the union is risk-averse, as shown in row (b). As explained above, this is caused by the employer best-responding to a more moderate competing offer by the union when the union is risk averse. As a result, when a risk-averse union does lose arbitration, the awarded wage increase is higher than when a risk-neutral union loses. In combination, these two aspects more than compensate the smaller wage increase obtained by the risk-averse union in the event that it wins arbitration, so the net result is an increase in expected awards.

Yet despite the larger arbitrated wage on average, Table 6, row (e) shows that the risk-averse union's certainty equivalent of arbitration is lower than in the risk neutral scenario because the risk premium of arbitration is sufficiently large. How do these effects of risk-aversion—the rise in the expected arbitrated wage increase and

the reduction in the union's certainty equivalent of arbitration—affect the relative strengths of the parties' positions in a dispute where settlement failure triggers arbitration? Intuitively—and also according to models of bargaining such as in Nash Jr (1950)—a party can extract a better outcome from bargaining as its prospects in the event of a disagreement improve. In settings where arbitration is the terminal dispute resolution procedure, arbitration serves as the disagreement outcome of bargaining. Table 6 shows that the union's risk aversion causes its certainty equivalent of arbitration to fall more than the employer's compared to the risk neutral baseline. Thus, somewhat paradoxically, risk aversion can weaken a party's position in a dispute where arbitration is the terminal procedure despite making it more likely to win the arbitration case.³⁴

7. CONCLUSION

We combine economic theory and empirics to study arbitration, a widely used method of resolving disputes. Our model of the three-way strategic interaction between two disputing parties and an arbitrator highlights the following features of arbitration: First, risk attitudes affect the strategic actions of the players and the outcomes that ensue; asymmetry in these risk attitudes can tilt outcomes in favor of one side or another. Second, arbitration is a game of communication with the arbitrator. Under final-offer arbitration, we establish identification of the model from the joint distribution of offers submitted by the disputing parties and the arbitration awards. Based on the identification strategy, we develop an estimator, which we then implement using data on wage arbitration between police and fire officer unions and their employers in the state of New Jersey. This is the first structural analysis of arbitration.

Our data affords us a rare opportunity to study in the field a cheap-talk and a non-cheap-talk version of a communication game—conventional and final-offer arbitration, respectively. Noting that the disputing parties' offers are further apart in conventional arbitration, we leverage our structural model to quantify the relative precision of information transmission in the cheap-talk game. We find that, in our application, the information communicated in conventional arbitration is less than half as precise as that in final-offer arbitration. However, the superior information in final-offer arbitration comes at the cost of constraining the arbitrator's choice of award to one

³⁴In Appendix F, we model the pre-arbitration interaction between the union and the employer to make explicit the connection between the parties' arbitration certainty equivalents and their relative bargaining positions in pre-arbitration negotiations.

of the parties' offers, so there is a trade-off between eliciting information and allowing more arbitrator discretion. On balance, we find that conventional arbitration achieves outcomes that are closer to the ideal outcome in our application.

When considering final-offer arbitration in isolation, we find that the more risk-averse party actually obtains superior outcomes (more favorable wages) on average, partly because it submits moderate offers that are more likely to be chosen by the arbitrator. Nonetheless, given the ex-ante uncertainty about the arbitration award, the risk-averse party ultimately has a lower certainty equivalent of arbitration than if it were risk neutral, which may weaken its position in a dispute where arbitration is the disagreement outcome.

Our analysis may be extended in various ways. Whereas we study one-dimensional information and actions in this paper, an important extension would be to characterize multidimensional disputes involving multidimensional information and action spaces. Another interesting question is to investigate more explicitly the possible dynamic linkages between arbitration cases. Finally, the questions we ask of arbitration have analogs in dispute resolution more generally. For example, the lack of discretion faced by arbitrators in final-offer arbitration is of a similar nature to the constraints that structured sentencing systems, such as sentencing guidelines and mandatory minimum sentences, pose on judges in criminal cases. Adapting our framework to the investigation of the trade-offs associated with judicial discretion, accounting for the possibility of strategic communication, would be an exciting avenue for further research.

REFERENCES

- Ambrus, Attila, Eric Chaney, and Igor Salitskiy**, "Pirates of the mediterranean: An empirical investigation of bargaining with asymmetric information," *Quantitative Economics*, 2018, *9* (1), 217–246.
- Ashenfelter, Orley**, "Arbitrator Behavior," *American Economic Review*, 1987, *77* (2), 342–346.
- **and David E. Bloom**, "Models of Arbitrator Behavior: Theory and Evidence," *American Economic Review*, 1984, *74* (1), 111–124.
- **and Gordon B. Dahl**, "Bargaining and the Role of Expert Agents: An Empirical Study of Final-Offer Arbitration," *Review of Economics and Statistics*, 2012, *94* (1), 116–132.

- , **Janet Currie, Henry S Farber, and Matthew Spiegel**, “An Experimental Comparison of Dispute Rates in Alternative Arbitration Systems,” *Econometrica*, 1992, *60* (6), 1407–1433.
- Austen-Smith, David and Jeffrey S. Banks**, “Cheap talk and burned money,” *Journal of Economic Theory*, 2000, *91* (1), 1–16.
- Babcock, Bruce A., E. Kwan Choi, and Eli Feinerman**, “Risk and probability premiums for CARA utility functions,” *Journal of Agricultural and Resource Economics*, 1993, pp. 17–24.
- Backus, Matthew, Thomas Blake, and Steven Tadelis**, “On the Empirical Content of Cheap-Talk Signaling: An Application to Bargaining,” *Journal of Political Economy*, 2019, *127* (4), 1599–1628.
- Bagwell, Kyle, Robert W. Staiger, and Ali Yurukoglu**, “Multilateral trade bargaining: A first look at the GATT bargaining records,” *American Economic Journal: Applied Economics*, 2020, *12* (3), 72–105.
- Barr, Michael S.**, “Mandatory arbitration in consumer finance and investor contracts,” *New York University Journal of Law and Business*, 2014, *11*, 793.
- Bazerman, Max H. and Henry S. Farber**, “Arbitrator Decision Making: When are Final Offers Important?,” *Industrial and Labor Relations Review*, 1985, *39* (1), 76–89.
- Bloom, David E.**, “Collective bargaining, compulsory arbitration, and salary settlements in the public sector: The case of New Jersey’s municipal police officers,” *Journal of Labor Research*, 1981, *2* (2), 369–384.
- , “Empirical Models of Arbitrator Behavior under Conventional Arbitration,” *The Review of Economics and Statistics*, 1986, *68* (4), 578–85.
- Carrell, Micael and Richard Bales**, “Considering final offer arbitration to resolve public sector impasses in times of concession bargaining,” *Ohio St. J. on Disp. Resol.*, 2013, *28*, 1.
- Carruth, Alan A. and Andrew J. Oswald**, “Miners’ wages in post-war Britain: An application of a model of trade union behaviour,” *The Economic Journal*, 1985, *95* (380), 1003–1020.
- Çelen, Boğaçhan and Onur Özgür**, “Final-offer arbitration with uncertainty averse parties,” *Games and Economic Behavior*, 2018, *109*, 484–500.
- Chan, David C., Matthew Gentzkow, and Chuan Yu**, “Selection with variation in diagnostic skill: Evidence from radiologists,” *The Quarterly Journal of Economics*, 2022, *137* (2), 729–783.

- Crawford, Gregory S. and Matthew Shum**, “Uncertainty and learning in pharmaceutical demand,” *Econometrica*, 2005, 73 (4), 1137–1173.
- Crawford, Vincent P.**, “On compulsory-arbitration schemes,” *Journal of Political Economy*, 1979, 87 (1), 131–159.
- , “A theory of disagreement in bargaining,” *Econometrica*, 1982, pp. 607–637.
- **and Joel Sobel**, “Strategic Information Transmission,” *Econometrica*, 1982, 50 (6), 1431–1451.
- Currie, Janet**, “Who Uses Interest Arbitration? The Case of British Columbia’s Teachers, 1947-1981.,” *Industrial and Labor Relations Review*, 1989, 42 (3), 363–79.
- DeGroot, Morris H.**, *Optimal Statistical Decisions*, Vol. 82, John Wiley & Sons, 2005.
- Dickinson, David L.**, “A comparison of conventional, final-offer, and “combined” arbitration for dispute resolution,” *ILR Review*, 2004, 57 (2), 288–301.
- Egan, Mark L., Gregor Matvos, and Amit Seru**, “Arbitration with uninformed consumers,” Technical Report, National Bureau of Economic Research 2018.
- Eraslan, Hülya KK.**, “Corporate bankruptcy reorganizations: estimates from a bargaining model,” *International Economic Review*, 2008, 49 (2), 659–681.
- Farber, Henry S.**, “Individual preferences and union wage determination: the case of the united mine workers,” *Journal of Political Economy*, 1978, 86 (5), 923–942.
- , “An Analysis of Final-Offer Arbitration,” *Journal of Conflict Resolution*, 1980, 24 (4), 683–705.
- **and Harry C. Katz**, “Interest arbitration, outcomes, and the incentive to bargain,” *ILR Review*, 1979, 33 (1), 55–63.
- **and Max H. Bazerman**, “The General Basis of Arbitrator Behavior: An Empirical Analysis of Conventional and Final-Offer Arbitration,” *Econometrica*, 1986, 54 (4), 819–844.
- Farmer, Amy and Paul Pecorino**, “Bargaining with informative offers: An analysis of final-offer arbitration,” *The Journal of Legal Studies*, 1998, 27 (2), 415–432.
- Feuille, Peter**, “Final Offer Arbitration and the Chilling Effect,” *Industrial Relations: A Journal of Economy and Society*, 1975, 14 (3), 302–310.
- Gibbons, Robert**, “Learning in Equilibrium Models of Arbitration,” *The American Economic Review*, 1988, 78 (5), 896–912.
- Haan, Thomas De, Theo Offerman, and Randolph Sloof**, “Money talks? An experimental investigation of cheap talk and burned money,” *International Economic Review*, 2015, 56 (4), 1385–1426.

- Jr, John F. Nash**, “The bargaining problem,” *Econometrica: Journal of the econometric society*, 1950, pp. 155–162.
- Kawai, Kei, Ken Onishi, and Kosuke Uetake**, “Signaling in Online Credit Markets,” *Journal of Political Economy*, 2022, 130 (6), 1364–1411.
- Larsen, Bradley J.**, “The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions,” *The Review of Economic Studies*, 2020.
- _____ and **Joachim Freyberger**, “How Well Does Bargaining Work in Consumer Markets? A Robust Bounds Approach,” Technical Report, National Bureau of Economic Research 2021.
- Lipsky, David B. and Ronald L. Seeber**, “The appropriate resolution of corporate disputes: A report on the growing use of ADR by US corporations,” Technical Report, Ithaca NY: Cornell Studies in Conflict and Dispute Resolution 1998.
- Madhavan, Ananth**, “Trading Mechanisms in Securities Markets,” *The Journal of Finance*, 1992, 47 (2), 607–641.
- Marburger, Daniel R. and John F. Scoggins**, “Risk and final offer arbitration usage rates: Evidence from major league baseball,” *Journal of Labor Research*, 1996, 17 (4), 735–745.
- Mas, Alexandre**, “Pay, reference points, and police performance,” *The Quarterly Journal of Economics*, 2006, 121 (3), 783–821.
- McCall, Brian P.**, “Interest arbitration and the incentive to bargain: A principal-agent approach,” *Journal of Conflict Resolution*, 1990, 34 (1), 151–167.
- Merlo, Antonio**, “Bargaining over governments in a stochastic environment,” *Journal of Political Economy*, 1997, 105 (1), 101–131.
- _____ and **Xun Tang**, “Identification and estimation of stochastic bargaining models,” *Econometrica*, 2012, 80 (4), 1563–1604.
- _____ and _____, “Bargaining with optimism: Identification and estimation of a model of medical malpractice litigation,” *International Economic Review*, 2019, 60 (3), 1029–1061.
- _____ and _____, “New results on the identification of stochastic bargaining models,” *Journal of Econometrics*, 2019, 209 (1), 79–93.
- Miller, Robert A.**, “Job matching and occupational choice,” *Journal of Political economy*, 1984, 92 (6), 1086–1120.
- Mnookin, Robert H.**, *Alternative dispute resolution*, Harvard Law School, 1998.
- Mylovanov, Tymofiy and Andriy Zapechelnyuk**, “Optimal arbitration,” *International Economic Review*, 2013, 54 (3), 769–785.
- Neale, Margaret A. and Max H. Bazerman**, “The role of perspective-taking

- ability in negotiating under different forms of arbitration,” *ilr Review*, 1983, *36* (3), 378–388.
- Olszewski, Wojciech**, “A welfare analysis of arbitration,” *American Economic Journal: Microeconomics*, 2011, *3* (1), 174–213.
- Pecorino, Paul, Michael Solomon, and Mark Van Boening**, “Bargaining with voluntary transmission of private information: An experimental analysis of final offer arbitration,” *Journal of Economic Behavior & Organization*, 2021, *191*, 334–366.
- Perry, Motty**, “An Example of Price Formation in Bilateral Situations: A Bargaining Model with Incomplete Information,” *Econometrica*, 1986, *54* (2), 313–321.
- Pratt, John W.**, “Risk Aversion in the Small and in the Large,” *Econometrica*, 1964, *32* (1), 122–136.
- Reilly, William J.**, “Distressed grievance procedures and their rehabilitation-discussion,” in “Labor Arbitration and Industrial Change, Proceedings of the Annual Meeting of the National Academy of Arbitrators” 1963, pp. 132–138.
- Sahni, Navdeep S. and Harikesh S. Nair**, “Does Advertising Serve as a Signal? Evidence from a Field Experiment in Mobile Search,” *The Review of Economic Studies*, 2020, *87* (3), 1529–1564.
- Samuelson, William F.**, “Final-offer arbitration under incomplete information,” *Management science*, 1991, *37* (10), 1234–1247.
- Schelling, Thomas C.**, “An essay on bargaining,” *The American Economic Review*, 1956, *46* (3), 281–306.
- Senate Standing Committee on Economics, Parliament of Australia**, “Treasury Laws Amendment (News Media and Digital Platforms Mandatory Bargaining Code) Bill 2020 (Bills Digest No. 48, 2020–21),” 2021.
- Sieg, Holger**, “Estimating a bargaining model with asymmetric information: Evidence from medical malpractice disputes,” *Journal of Political Economy*, 2000, *108* (5), 1006–1021.
- Silveira, Bernardo S.**, “Bargaining with asymmetric information: An empirical study of plea negotiations,” *Econometrica*, 2017, *85* (2), 419–452.
- Silver-Greenberg, Jessica and Robert Gebeloff**, “Arbitration Everywhere, Stacking the Deck of Justice,” *New York Times*, October 2015. <https://www.nytimes.com/2015/11/01/business/dealbook/arbitration-everywhere-stacking-the-deck-of-justice.html>.
- Slater, Joseph**, “Interest arbitration as alternative dispute resolution: The history from 1919 to 2011,” *Ohio St. J. on Dispute Resolution*, 2013, *28* (2), 387–418.

- Spence, Michael**, “Job Market Signaling,” *Quarterly Journal of Economics*, 1973, 87 (3), 355–374.
- Stevens, Carl M.**, “Is Compulsory Arbitration Compatible With Bargaining?,” *Industrial Relations: A Journal of Economy and Society*, 1966, 5 (2), 38–52.
- Stokes, Greg**, “Solomon’s Wisdom: An Early Analysis of the Effects of the Police and Fire Interest Arbitration Reform Act in New Jersey,” *Journal of Collective Negotiations in the Public Sector*, 1999, 28 (3), 219–232.
- Sweeting, Andrew, James W. Roberts, and Chris Gedge**, “A model of dynamic limit pricing with an application to the airline industry,” *Journal of Political Economy*, 2020, 128 (3), 1148–1193.
- Waldfogel, Joel**, “The selection hypothesis and the relationship between trial and plaintiff victory,” *Journal of Political Economy*, 1995, 103 (2), 229–260.
- Watanabe, Yasutora**, “Learning and bargaining in dispute resolution: Theory and evidence from medical malpractice litigation,” *Unpublished manuscript*, 2006.

APPENDICES FOR ONLINE PUBLICATION


APPENDIX A. SUPPLEMENTARY TABLES AND FIGURES

TABLE A1. Offer Aggressiveness and Employer Win Probability, 1978-1995

	(1)	(2)	(3)
Union final offer residual	0.218 (0.043)		0.140 (0.049)
Employer final offer residual		0.242 (0.046)	0.169 (0.052)
Constant	-0.324 (0.054)	-0.334 (0.054)	-0.333 (0.054)
Observations	579	579	579

Notes: Table reports Probit results. The unit of observation is a case. In all specifications, the sample consists of cases from the ARB_F data set, which are resolved by final-offer arbitration. The dependent variable is a dummy indicating whether the employer wins the arbitration. The regressors are residuals of regressions of the final offers by the union and the employer on all the covariates in column (1) of Table 2. Standard errors provided in parentheses.


FIGURE A1. Scatter Plot of Final Offers, 1978–1995



Notes: Employer and union final offers in all cases from the ARB_F data set. The 45 degree line is marked with a dotted line.

FIGURE A2. Histograms of Arbitration Data

(a) Midpoint of Union and Employer FOA Offers, 1978–1995



(b) CA Arbitrated Wages, 1996–2000


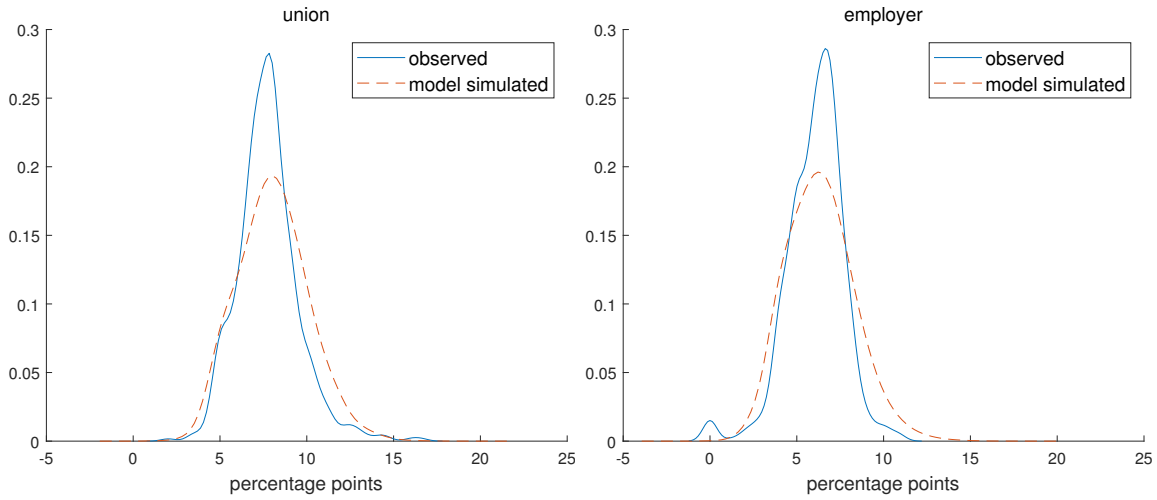



FIGURE A3. Model Fit: Final Offers, 1978-1995



Notes: Figures display kernel density of observed vs. model-simulated final offers by the union and the employer, respectively.

FIGURE A4. Count of arbitration awards by year of arbitration



Notes: The policy change from FOA to CA occurred in 1996.

APPENDIX B. FINAL OFFER AND CONVENTIONAL ARBITRATION:
SUPPLEMENTARY EVIDENCE

As a complement to the counterfactual analysis presented in Section 6.1, this Appendix compares the arbitration outcomes in final-offer (FOA) and conventional arbitration (CA) using a descriptive regression exercise. Recall that, in our setting, FOA was the default dispute resolution method until 1995, whereas, from 1996 onward, cases were resolved by CA. We exploit this institutional change in the following specification:

$$Outcome_i = \mu_0 + \mu_1 Conventional_i + \mu_2 X_i + \iota_i, \quad (A.1)$$

where the unit of observation is a case, denoted by i , and ι_i is an error term. As the dependent variable, $Outcome_i$, we consider the analogs of Table 4 outcomes, namely: (i) the difference between the offers made by the union and the employer; (ii) the difference between the wage increase decided by the arbitrator and the midpoint of the offers made by the parties; and (iii) the arbitrated wage increase. The regressor of interest is $Conventional_i$, a dummy that indicates whether case i is decided after 1996—that is, by CA. The vector X_i contains all of the covariates included in column (1) of Table 2 in the main text, except for the year-group fixed effects. Instead of controlling for year groups, we estimate (A.1) using only data on cases resolved from 1993 onward, so the FOA data used in the regression analysis constitutes only the last year group from the estimation sample employed in the main text (see Section 2.3 for information on the year-group fixed effects).

Table A2 presents OLS estimates of (A.1). Relative to FOA, CA is associated with a wider gap between the offers made by the union and the employer, as shown in column (1). Column (2) shows that, taking the midpoint between the parties' offers as a reference, the awards chosen by the arbitrator are smaller in CA than in FOA.¹ Column (3) shows that CA cases are associated with a lower absolute arbitrated wage increase than are FOA cases. These findings mirror our results from Section 6.1, albeit with Section 6.1 estimating an effect of smaller magnitude for column (3).

It is worth stressing that, besides the obvious methodological distinctions, the regression presented in this Appendix and the counterfactual analysis in Section 6.1 are based on different samples. The latter provides a comparison between *observed*

¹In FOA, the award tends to be above the offer midpoint because the union wins arbitration more often than the employer. In CA, the concept of one party winning or losing does not apply. But we can assess whether, in expectation, the award is closer to the union's offer or to that by the employer, and by how much. This is the purpose of comparing the award to the midpoint of offers, as we do in Table A2, column (2).

CA cases post-1996 and FOA outcomes that are *simulated*, given the covariates of the *same* post-1996 cases. In contrast, the regressions presented here compare only observed cases—using 1993-1995 data on FOA cases and 1996-2000 data on CA cases, so that covariates of the cases would be different for the two arbitration designs. Thus, “differences” between results of the two analyses need not imply a contradiction. In Table A2, we think column (3) would be the most influenced by these differences in covariate samples, while columns (1) and (2) would be more robust because they examine outcomes that constitute within-case differences. Indeed, Section 6.1, by enabling a comparison of arbitration designs given the *same* set of covariates, informs us that the estimated OLS coefficient in Table A2, column (3) may be exaggerated in magnitude.

Overall, the results from the reduced-form and structural approaches corroborate each other here and provide further credibility to the subsequent analyses in the main text that are motivated by these comparisons.

TABLE A2. FOA vs. CA: Offers and Case Outcomes (1993-2000)

	(1)	(2)	(3)
	Difference between Offers	Arb. Wage - Offer Midpoint	Arbitrated Wage Increase
Conventional	1.832 (0.319)	-0.357 (0.188)	-0.802 (0.163)
Observations	158	158	158
R^2	0.394	0.175	0.416
Adjusted R^2	0.280	0.019	0.305

Controls: number of years covered by the contract; 12-month percent change in the CPI; *othermuni* (see Section 2.3 in main text for details); log of taxable property per capita; quantile rank of median household income among NJ municipalities; log of population; population density; a dummy indicating a contract for fire officers; a dummy indicating that the employer is a county; and the credit rating assigned to municipal debt obligations by Moody’s Investors’ Service.

Notes: Table reports OLS results. The unit of observation is a case. In all specifications, the sample consists of cases decided by final-offer arbitration (ARB_F data) from 1993-1995 and cases resolved by conventional arbitration (ARB_C data) from 1996-2000. The regressor of interest is a dummy indicating whether the case was decided by conventional arbitration. Standard errors provided in parentheses.

APPENDIX C. PROOFS

Proof of Proposition 1.

Proof. We adopt a “guess and verify” approach for the proof. Assume that offers take the form $y_u(s_p) = M_p(s_p) + \delta_u$ and $y_e(s_p) = M_p(s_p) - \delta_e$, where δ_u and δ_e do not depend on s_p .

First, we characterize the arbitrator’s inference and the decision rule that best responds to the supposed $y_u(s_p)$, $y_e(s_p)$. As derived in the text following Proposition 1, the arbitrator’s best response given the supposed $y_u(s_p)$, $y_e(s_p)$ is to infer s_p by the inference rule

$$s_p(\bar{y}) = \frac{(h + h_\varepsilon) [\bar{y} + (\delta_e - \delta_u)/2] - hm}{h_\varepsilon}.$$

Also, as derived in the text, the arbitrator then chooses y_e if and only if

$$s_a < \frac{h_\varepsilon \bar{y} + h(\bar{y} - m) + h_\varepsilon (\bar{y} - s_p(\bar{y}))}{h_\varepsilon} = \bar{y} - \left(\frac{h + h_\varepsilon}{h_\varepsilon} \right) \frac{\delta_e - \delta_u}{2} \equiv S(\bar{y}).$$

Second, we confirm that there exists a unique pair δ_u , δ_e such that the final offer strategies $y_u(s_p) = M_p(s_p) + \delta_u$ and $y_e(s_p) = M_p(s_p) - \delta_e$ in turn best respond to the inference and decision rules above and to one another. By Assumption 1, the parties’ belief about the distribution of s_a conditional on s_p is normal with mean $M_p(s_p)$ and precision $H = [h_\varepsilon(h + h_\varepsilon)] / (h + 2h_\varepsilon)$. Let $\Phi(\cdot)$ and $\phi(\cdot)$ be the standard normal cumulative distribution and density functions, respectively. Then the decision rule above implies that the arbitrator selects y_e with probability $\Phi([S(\bar{y}) - M_p(s_p)]\sqrt{H})$.

We can then rewrite the problems solved by the union and the employer, respectively, as

$$\begin{aligned} & \max_{\delta_u} u_u(M_p(s_p) - \delta_e) \Phi([S(\bar{y}) - M_p(s_p)]\sqrt{H}) \\ & \quad + u_u(M_p(s_p) + \delta_u) \left[1 - \Phi([S(\bar{y}) - M_p(s_p)]\sqrt{H}) \right], \\ \text{and } & \max_{\delta_e} u_e(M_p(s_p) - \delta_e) \Phi([S(\bar{y}) - M_p(s_p)]\sqrt{H}) \\ & \quad + u_e(M_p(s_p) + \delta_u) \left[1 - \Phi([S(\bar{y}) - M_p(s_p)]\sqrt{H}) \right]. \end{aligned}$$

The corresponding first-order conditions are

$$\frac{\sqrt{H}}{2} \frac{\phi([S(\bar{y}) - M_p(s_p)]\sqrt{H})}{1 - \Phi([S(\bar{y}) - M_p(s_p)]\sqrt{H})} = \frac{\rho}{\exp(\rho(\delta_u + \delta_e)) - 1},$$

$$\text{and } \frac{\sqrt{H}}{2} \frac{\phi([S(\bar{y}) - M_p(s_p)]\sqrt{H})}{\Phi([S(\bar{y}) - M_p(s_p)]\sqrt{H})} = \frac{1}{\delta_u + \delta_e},$$

where we use the fact that the derivative of $S(\bar{y})$ with respect to the union's choice of δ_u and the employer's choice of δ_e are $1/2$ and $-1/2$, respectively.

In equilibrium, δ_u and δ_e must satisfy these FOCs with $M_p(s_p) = (\bar{y} + (\delta_e - \delta_u)/2)$. Plugging in this expression and rearranging, we find that the equilibrium δ_u and δ_e must satisfy

$$\begin{aligned} \frac{\sqrt{H}}{2} \frac{\phi(\eta(\delta_u - \delta_e)/2)}{1 - \Phi(\eta(\delta_u - \delta_e)/2)} &= \frac{\rho}{\exp(\rho(\delta_u + \delta_e)) - 1}, \\ \text{and } \frac{\sqrt{H}}{2} \frac{\phi(\eta(\delta_u - \delta_e)/2)}{\Phi(\eta(\delta_u - \delta_e)/2)} &= \frac{1}{\delta_u + \delta_e}, \end{aligned}$$

where $\eta \equiv \sqrt{H}(h + 2h_e)/h_e$. These correspond to (4) and (5) in the text.

To show that there exists a unique pair δ_u, δ_e that solves the system of equations implied by these first-order conditions, define shorthand $t \equiv \eta(\delta_u - \delta_e)/2$, $d_1 \equiv \delta_u + \delta_e$, $f(d_1) \equiv \rho/(\exp(\rho d_1) - 1)$, $\lambda \equiv \phi/(1 - \Phi)$ and $\tilde{\lambda} \equiv \phi/\Phi$. We can rewrite (4) and (5) as

$$\frac{\sqrt{H}}{2} \lambda(t) = f(d_1) \quad \text{and} \quad \frac{\sqrt{H}}{2} \tilde{\lambda}(t) = 1/d_1. \quad (\text{A.2})$$

This system admits a solution in $t \in \mathbb{R}$ and $d_1 \in \mathbb{R}_+$ if and only if

$$\frac{\sqrt{H}}{2} \lambda(t) = f\left(\frac{2}{\sqrt{H}\tilde{\lambda}(t)}\right) \quad (\text{A.3})$$

admits a solution in $t \in \mathbb{R}$. By construct, λ is increasing, while $\tilde{\lambda}$ and f are decreasing in t and d_1 , respectively. As $t \rightarrow -\infty$, we know that $\lambda(t) \rightarrow 0$, $\tilde{\lambda}(t) \rightarrow \infty$, and the r.h.s of (A.3) diverges to ∞ . On the other hand, as $t \rightarrow \infty$, we have that $\lambda(t) \rightarrow \infty$, $\tilde{\lambda}(t) \rightarrow 0$, and the r.h.s. of (A.3) converges to 0. Therefore both sides of (A.3) are strictly monotonic in different directions, implying existence of a unique solution in t . Given t , (A.2) pins down a unique d_1 . Then, since t determines the difference between δ_u and δ_e and d_1 determines their sum, existence and uniqueness of t and d_1 yields existence and uniqueness of the values of δ_u and δ_e that satisfy (4) and (5).

Finally, as s_p is absent from (4) and (5), we verify that neither δ_u nor δ_e vary with the parties' signal s_p . \square

Proof of Proposition 2.

Proof. (i) Let $d_1 \equiv \delta_u + \delta_e$, the distance between final offers. In a proof by contradiction, suppose $h' > h$ and $d_1(h') \geq d_1(h)$. As the right-hand sides of (A.2) both

decrease in d_1 , we have $\sqrt{H(h')}\lambda(t(h')) \leq \sqrt{H(h)}\lambda(t(h))$ and $\sqrt{H(h')}\tilde{\lambda}(t(h')) \leq \sqrt{H(h)}\tilde{\lambda}(t(h))$. Since H is strictly increasing in h , this is only possible if $\lambda(t(h')) < \lambda(t(h))$ and $\tilde{\lambda}(t(h')) < \tilde{\lambda}(t(h))$. However, by definition, $\lambda(\cdot)$ is strictly increasing, while $\tilde{\lambda}(\cdot)$ is strictly decreasing, so it is impossible for these two inequalities to be satisfied simultaneously. Therefore, $d_1(h') < d_1(h)$ by contradiction. Repeat the same proof replacing h with h_ε to show that d_1 is strictly decreasing in h_ε .

(ii) While we use risk-neutrality for the employer and CARA utility for the union throughout this paper, here we relax the employer's risk-neutrality to prove a more general point. Let $U_u(\cdot)$ and $U_e(\cdot)$ be notation for the parties' CARA utility functions, which may differ in their risk aversion parameters. Taking a ratio of (4) and (5) yields

$$\frac{\Phi(\eta(\delta_u - \delta_e)/2)}{1 - \Phi(\eta(\delta_u - \delta_e)/2)} = \left(\frac{U_e(-y_e) - U_e(-y_u)}{U_u(y_u) - U_u(y_e)} \right) \frac{U'_u(y_u)}{U'_e(-y_e)}. \quad (\text{A.4})$$

Now define a function $\tilde{U}_e(\cdot)$ such that $\tilde{U}_e(z + (y_u + y_e)) \equiv U_e(z)$. Note that, in terms of absolute risk aversion, if $U_u(\cdot)$ is more (less) risk-averse than $U_e(\cdot)$, it is also more (less) risk-averse than $\tilde{U}_e(\cdot)$. We can rewrite the equation above as

$$\frac{\Phi(\eta(\delta_u - \delta_e)/2)}{1 - \Phi(\eta(\delta_u - \delta_e)/2)} = \left(\frac{\tilde{U}_e(y_u) - \tilde{U}_e(y_e)}{U_u(y_u) - U_u(y_e)} \right) \frac{U'_u(y_u)}{\tilde{U}'_e(y_u)}.$$

By equation (22) in Pratt (1964), the r.h.s. of the above equation is < 1 if the union is more risk-averse, $= 1$ if the parties are equally risk-averse, and > 1 if the employer is more risk-averse. Then by the l.h.s. of the equation and properties of the standard normal cdf $\Phi(\cdot)$, $\delta_u < \delta_e$ if the union is more risk-averse, $\delta_u = \delta_e$ if the parties are equally risk-averse, and $\delta_u > \delta_e$ if the employer is more risk-averse.

Meanwhile, the l.h.s. above is the odds of the employer winning, by definition. Thus, the more risk-averse party wins more often in expectation. This proof is closely related to that of Farber (1980). \square

Proof of Proposition 3.

Proof. Denote the final offers by the union and the employer, respectively, by $y_u(s_p, h_\varepsilon)$ and $y_e(s_p, h_\varepsilon)$. From Proposition 1, we have $y_u(s_p, h_\varepsilon) = M_p(s_p, h_\varepsilon) + \delta_u(h_\varepsilon)$ and $y_e(s_p, h_\varepsilon) = M_p(s_p, h_\varepsilon) - \delta_e(h_\varepsilon)$. Define $d_1(h_\varepsilon) \equiv y_u(s_p, h_\varepsilon) - y_e(s_p, h_\varepsilon) = \delta_u(h_\varepsilon) + \delta_e(h_\varepsilon)$ and $d_2(h_\varepsilon) \equiv (\delta_u(h_\varepsilon) - \delta_e(h_\varepsilon))$. Also, by (6), in equilibrium the arbitrator chooses the employer's final offer with probability $\Phi(\eta(h_\varepsilon)(\delta_u(h_\varepsilon) - \delta_e(h_\varepsilon))/2)$, where $\eta(h_\varepsilon) \equiv \sqrt{H(h_\varepsilon)}(h + 2h_\varepsilon)/h_\varepsilon$ and $H(h_\varepsilon) \equiv h_\varepsilon(h + h_\varepsilon)/(h + 2h_\varepsilon)$.

First, we show that ρ is identified. From (7), we have

$$\frac{\Phi(\eta(h_\varepsilon)d_2(h_\varepsilon)/2)}{1 - \Phi(\eta(h_\varepsilon)d_2(h_\varepsilon)/2)} = \frac{\rho d_1(h_\varepsilon)}{\exp(\rho d_1(h_\varepsilon)) - 1}.$$

Let $odds(y_u - y_e)$ denote the observed odds that the employer's final offer is chosen by the arbitrator, conditional on the observed offer difference $y_u - y_e$. Proposition 2(i) shows that $d_1(h_\varepsilon)$ is strictly decreasing in h_ε , allowing us to use $h_\varepsilon = d_1^{-1}(y_u - y_e)$ and write

$$odds(y_u - y_e) = \frac{\Phi(\eta(d_1^{-1}(y_u - y_e))d_2(d_1^{-1}(y_u - y_e))/2)}{1 - \Phi(\eta(d_1^{-1}(y_u - y_e))d_2(d_1^{-1}(y_u - y_e))/2)}. \quad (\text{A.5})$$

Together, the equations above imply

$$odds(y_u - y_e) = \frac{\rho(y_u - y_e)}{\exp(\rho(y_u - y_e)) - 1}. \quad (\text{A.6})$$

From Theorem 1 and equation (22) in Pratt (1964), the r.h.s. is strictly decreasing in ρ , so the equation above identifies this parameter.

Next, we show the identification of h and $G_{h_\varepsilon}(\cdot)$. First, since $\Phi(x)/[1 - \Phi(x)]$ is strictly increasing in x , (A.5) identifies the product $\eta(d_1^{-1}(y_u - y_e))d_2(d_1^{-1}(y_u - y_e))$. Plugging this value into the left-hand side of (4) then identifies $H(d_1^{-1}(y_u - y_e))$, as the r.h.s. of that equation is a ratio of two identified terms. Rearranging the definition of $H(h_\varepsilon)$ gives

$$\frac{1}{H(h_\varepsilon)} = \frac{1}{h_\varepsilon} + \frac{1}{h + h_\varepsilon} = \frac{h}{h_\varepsilon} \left(\frac{1}{h} + \frac{1}{h} \frac{1}{1 + \frac{h}{h_\varepsilon}} \right). \quad (\text{A.7})$$

Meanwhile, from the definition of $M_p(s_p, h_\varepsilon)$, we have that

$$\text{Var}[M_p(s_p, h_\varepsilon) | h_\varepsilon] = \left(\frac{h_\varepsilon}{h + h_\varepsilon} \right)^2 \text{Var}[s_p | h_\varepsilon] = \frac{1}{h} \left(\frac{1}{1 + \frac{h}{h_\varepsilon}} \right), \quad (\text{A.8})$$

where the l.h.s. is an observed quantity because

$$\begin{aligned} \text{Var}[M_p(s_p, h_\varepsilon) | h_\varepsilon = d_1^{-1}(y_u - y_e)] &= \text{Var}[y_u(s_p, h_\varepsilon) - \delta_u(h_\varepsilon) | h_\varepsilon = d_1^{-1}(y_u - y_e)] \\ &= \text{Var}[y_u(s_p, h_\varepsilon) | h_\varepsilon = d_1^{-1}(y_u - y_e)] \\ &= \text{Var}[y_u | y_u - y_e]. \end{aligned}$$

Equations (A.14) and (A.8) thus form a system of equations that can be solved for h and h_ε . Specifically, we rearrange (A.8) as

$$\frac{h}{h_\varepsilon} = \frac{1}{h \text{Var}[y_u | y_u - y_e]} - 1.$$

Plugging this into (A.7) gives

$$\frac{1}{H(d_1^{-1}(y_u - y_e))} = \left(\frac{1}{h \text{Var}[y_u | y_u - y_e]} - 1 \right) \left(\frac{1}{h} + \text{Var}[y_u | y_u - y_e] \right),$$

which corresponds to (8) in the text. The only unknown in the equation above is h , and the right-hand side is strictly decreasing in this parameter. Hence, this equation identifies h , which, in turn, identifies h_ε by (A.8). As the distribution of $y_u - y_e$ is observed, and we identify $h_\varepsilon = d_1^{-1}(y_u - y_e)$ for any value of $y_u - y_e$, we have nonparametric identification of $G_{h_\varepsilon}(\cdot)$.

Identification of h and h_ε implies identification of $\eta(h_\varepsilon)$. Then $d_2(h_\varepsilon)$ is identified since the product $\eta(h_\varepsilon)d_2(h_\varepsilon)$ is known. So we know both $d_2(h_\varepsilon)$ and $d_1(h_\varepsilon)$, implying recovery of $\delta_u(h_\varepsilon)$ and $\delta_e(h_\varepsilon)$ for all h_ε in the support of $G_{h_\varepsilon}(\cdot)$.

Finally, we identify the parameter m . We have

$$\mathbb{E}[M_p(s_p, h_\varepsilon)] = \mathbb{E}[\mathbb{E}[M_p(s_p, h_\varepsilon) | h_\varepsilon]] = \mathbb{E}\left[\frac{hm + h_\varepsilon \mathbb{E}[s_p | h_\varepsilon]}{h + h_\varepsilon}\right] = m.$$

Therefore, we have

$$\begin{aligned} m &= \mathbb{E}[\mathbb{E}[M_p(s_p, h_\varepsilon) | h_\varepsilon]] \\ &= \mathbb{E}[\mathbb{E}[y_u - \delta_u(h_\varepsilon) | h_\varepsilon]], \end{aligned}$$

where the right-hand side is now known. \square

Identifying the Employer's Risk Attitude. Suppose we allow CARA utility for both the union and the employer, so that ρ_u and ρ_e are the union's and employer's CARA parameters, respectively. By equation (A.4), the odds of the employer winning case i in equilibrium equals

$$\frac{\exp(\rho_e d_{1i}) - 1}{\exp(\rho_u d_{1i}) - 1} \frac{\rho_u}{\rho_e},$$

where d_{1i} is the difference between union and employer final offers in case i . Given variation in d_{1i} , the expression above yields many identifying equations, allowing estimation of both ρ_u and ρ_e as long as $\rho_u \neq \rho_e$. Estimating ρ_u and ρ_e using a minimum distance estimator based on the above, we obtain $\hat{\rho}_e \approx 0$.

APPENDIX D. ASYMMETRIC SIGNAL PRECISION

In this section, we extend the model in Section 3 to allow the arbitrator's signal precision to be *different* from that of the parties. Formally, we relax part (iii) of Assumption 1 as follows:

ASSUMPTION A1. *Part (i) and (ii) in Assumption 1 hold; (iii) the distributions of ε_a and ε_p are normal with mean zero and precision h_a and h_p , respectively.*

All other model elements remain the same as in Section 3. We characterize the equilibrium of this new model with asymmetric signal precision, prove identification of the model primitives and estimate them. We collect relevant proofs in Section D.4, focusing on differences from the original proofs to avoid repetition of text.

D.1. Equilibrium. When it comes to the model equilibrium and its properties summarized by Propositions 1 and 2, we find that the same conclusions hold as before, with h_p and h_a appropriately replacing h_e . Defining $\tilde{M}_p(s_p) \equiv (hm + h_p s_p)/(h + h_p)$, we restate the analogous propositions in terms of h_p and h_a below.

PROPOSITION A1. *Under Assumption A1, there exists a separating Perfect Bayesian Equilibrium of the arbitration game in which the final offers by the union and the employer have the form $y_u(s_p) = \tilde{M}_p(s_p) + \tilde{\delta}_u$ and $y_e(s_p) = \tilde{M}_p(s_p) - \tilde{\delta}_e$. The terms $\tilde{\delta}_u$ and $\tilde{\delta}_e$ are unique and do not depend on the signal s_p .*

PROPOSITION A2. *The equilibrium characterized in Proposition A1 is such that: (i) the distance between final offers ($\tilde{\delta}_u + \tilde{\delta}_e$) is strictly decreasing in the precision parameters h , h_p and h_a ; and (ii) the more risk-averse party chooses a final offer that is less distant from $\tilde{M}_p(s_p)$ —i.e., a smaller $\tilde{\delta}$ —and wins more often in expectation.*

D.2. Identification and Estimation.

D.2.1. Identification. Consider the data-generating process described in Section 4.1 with the following modifications. First, we treat the precision of the signals received by the parties, $h_{p,i}$, as a random variable drawn independently across cases i from a distribution $G_{h_p}(\cdot)$. Second, we let $h_{a,i} = \mu h_{p,i}$ for a constant $\mu > 0$ that represents how much more/less precise the arbitrator's signal is relative to that of the parties.

In this model, the primitives include: the union's risk aversion parameter, ρ ; the parameters of the fair wage increase distribution, m and h ; the distribution of the parties' signal precision, $G_{h_p}(\cdot)$; and the arbitrator-to-parties precision ratio, μ .

PROPOSITION A3. *Under Assumption A1 and the equilibrium of Proposition A1, the model primitives ρ , m , h , μ and the distribution $G_{h_p}(\cdot)$ are identified from the joint distribution of final offers (y_u, y_e) and the arbitrator's decision y .*

The intuition for identifying the model parameters other than the newly added μ is analogous to that presented in Section 4.2. As for μ , the arbitrator's signal precision ratio, it is identified separately from the parties' signal precision h_p because these have distinct effects on the final offer distribution. For simple intuition, consider the case of symmetric risk attitudes, given which the offer midpoint is exactly equal to $\tilde{M}_p(s_p)$. It is apparent that h_p affects both the variance of offer midpoints (because $\tilde{M}_p(s_p)$ is a function of h_p) and the distance between union and employer offers (by Proposition A2), while μ decreases the latter without affecting the former.

D.2.2. Estimation Procedure. Revisiting the key identifying equations from Section 5, equation (9) becomes

$$\tilde{H}_i \equiv \frac{\mu h_{p,i} [h_i + h_{p,i}]}{h_i + (1 + \mu)h_{p,i}}. \quad (\text{A.9})$$

Equation (11) now becomes

$$0 = \left[V_i \left(\frac{1}{\tilde{H}_i} + V_i \right) \right]^{\frac{1}{2}} - \left[\frac{1}{\mu h_i^2} + \left(1 - \frac{1}{\mu}\right) \frac{V_i}{h_i} \right]^{\frac{1}{2}}, \quad (\text{A.10})$$

where V_i is shorthand for $\text{Var}(y_{u,i} | \tilde{d}_{1,i}, x_i)$ and $\tilde{d}_{1,i} \equiv y_{u,i} - y_{e,i}$. When $\mu = 1$, (A.10) simplifies to (11) because $[h_i^{-2} + (1 - 1)V_i/h_i]^{\frac{1}{2}} = 1/h_i$. Equation (12) now becomes

$$\tilde{d}_{2,i}(\mu) = \frac{\mu h_{p,i} \Phi^{-1}(p_i)}{\sqrt{\tilde{H}_i} [h_i + (1 + \mu)h_{p,i}]} = \frac{\Phi^{-1}(p_i)}{\sqrt{\tilde{H}_i}} / \left[1 + \frac{1}{\mu} \left(\frac{1}{h_i V_i} \right) \right], \quad (\text{A.11})$$

where the second equality uses $h_i/h_{p,i} + 1 = 1/(h_i V_i)$ (see equation (A.15)).

We estimate ρ using the same steps as in Section 5. For estimating m and h , we maintain the specifications $m(x_i; \theta_m) = x_i \theta_m$ and $h(x_i; \theta_h) = 1/\exp(x_i \theta_h)$. Then let $\zeta_{1,i}(\theta_h; \mu)$ refer to the right-hand side of (A.10) evaluated at $h_i = h(x_i; \theta_h)$. Now, for any given value of μ , we can estimate the remaining model parameters as follows. Based on (A.10), we first estimate $\theta_h(\mu)$ as

$$\hat{\theta}_h(\mu) \equiv \arg \min_{\theta_h} \sum_i \zeta_{1,i}(\theta_h; \mu)^2.$$

Let $\hat{\zeta}_{1,i}(\mu) \equiv \zeta_{1,i}(\hat{\theta}_h(\mu); \mu)$. Second, we estimate $h_{p,i}(\mu)$ for each arbitration case in the sample by solving for $h_{p,i}$ in (A.9). Third, defining

$$\zeta_{2,i}(\theta_m; \mu) \equiv \frac{y_{u,i} + y_{e,i}}{2} - \hat{d}_{2,i}(\mu) - m(x_i; \theta_m),$$

we estimate $\theta_m(\mu)$ as

$$\hat{\theta}_m(\mu) \equiv \arg \min_{\theta_m} \sum_i \zeta_{2,i}(\theta_m; \mu)^2.$$

Let $\hat{\zeta}_{2,i}(\mu) \equiv \zeta_{2,i}(\hat{\theta}_m(\mu); \mu)$. Finally, we estimate μ by minimizing the criterion

$$\hat{\mu} \equiv \arg \min_{\mu} \sum_i w_1 \left[\hat{\zeta}_{1,i}(\mu) \right]^2 + w_2 \left[\hat{\zeta}_{2,i}(\mu) \right]^2.$$

The weights w_1 and w_2 are the inverse of the empirical variance of $\zeta_{1,i}(\hat{\theta}_{h,0}; \mu_0)$ and $\zeta_{2,i}(\hat{\theta}_{m,0}; \mu_0)$, respectively, where μ_0 refers to the value of μ originally used in the main text, i.e., $\mu_0 = 1$, and $\hat{\theta}_{h,0}$, $\hat{\theta}_{m,0}$ refer to the original estimates reported in the main text given $\mu_0 = 1$. These serve the role of “first-step estimates” allowing us to obtain the weights w_1 and w_2 .

D.3. Estimated signal precision ratio μ . Using the estimation procedure described in Section D.2.2, we obtain $\hat{\mu} = 0.90$ as the estimated ratio of arbitrator’s signal precision $h_{a,i}$ to parties’ signal precision $h_{p,i}$. The 95% bootstrap confidence interval of $\hat{\mu}$ based on 200 bootstrap samples is [0.55, 2.30].

D.4. Proofs for Propositions in Appendix D.1 and D.2.

Proof of Proposition A1. We adopt a “guess-and-verify” approach as in the proof of Proposition 1, assuming the parties’ final-offer strategies take the form $y_u(s_p) = \tilde{M}_p(s_p) + \tilde{\delta}_u$ and $y_e(s_p) = \tilde{M}_p(s_p) - \tilde{\delta}_e$, where $\tilde{\delta}_u$ and $\tilde{\delta}_e$ do not depend on s_p . The proof follows the exact steps of the proof of Proposition 1 with the following new definitions and notation. The arbitrator’s inference rule in equation (2) now becomes

$$\tilde{s}_p(\bar{y}) = \frac{(h + h_p) \left[\bar{y} + (\tilde{\delta}_e - \tilde{\delta}_u)/2 \right] - hm}{h_p}. \quad (\text{A.12})$$

The arbitrator’s updated expectation of the ideal wage increase given s_a and $\tilde{s}_p(\bar{y})$ becomes

$$\tilde{y}_a(s_a, y_u, y_e) = \frac{hm + h_p \tilde{s}_p(\bar{y}) + h_a s_a}{h + h_p + h_a}.$$

The arbitrator's decision rule for selecting y_e (equation (3)) now becomes

$$s_a < \frac{h_a \bar{y} + h(\bar{y} - m) + h_p(\bar{y} - \tilde{s}_p(\bar{y}))}{h_a} = \bar{y} - \left(\frac{h + h_p}{h_a} \right) \frac{\tilde{\delta}_e - \tilde{\delta}_u}{2} \equiv \tilde{S}(\bar{y}). \quad (\text{A.13})$$

Then the probability that the arbitrator selects y_e is $\Phi([\tilde{S}(\bar{y}) - \tilde{M}_p(s_p)]\sqrt{\tilde{H}})$, with $\tilde{H} \equiv (h_a(h + h_p))/(h + h_p + h_a)$. The rest of the proof follows the same argument as the proof of Proposition 1, with $M_p(s_p), \delta_e, \delta_u, H$ replaced by $\tilde{M}_p(s_p), \tilde{\delta}_e, \tilde{\delta}_u, \tilde{H}$ respectively, and with $\tilde{\eta} \equiv \sqrt{\tilde{H}}(h + h_p + h_a)/h_a$. \square

Proof of Proposition A2. Proof of (i) is almost identical to that in Proposition 2. The only difference is that the last step now uses the fact that \tilde{H} is increasing in both h_a and h_p , as well as in h . Proof of (ii) is identical to that in Proposition 2. \square

Proof of Proposition A3. Denote the final offers by the union and the employer by $y_u(s_p, h_p)$ and $y_e(s_p, h_p)$ respectively. Note that we now write h_p as an explicit argument in the final-offer strategies.

The first step is identify ρ , the risk aversion parameter in the union's utility. The argument for recovering ρ is almost identical to that in Proposition 3, only with $\delta_u(h_\varepsilon), \delta_e(h_\varepsilon), d_1(h_\varepsilon), d_2(h_\varepsilon)$ now replaced by $\tilde{\delta}_u(h_p), \tilde{\delta}_e(h_p), d_1(h_p) \equiv y_u(s_p, h_p) - y_e(s_p, h_p) = \tilde{\delta}_u(h_p) + \tilde{\delta}_e(h_p), d_2(h_p) \equiv \tilde{\delta}_u(h_p) - \tilde{\delta}_e(h_p)$ respectively, and with $\eta(h_\varepsilon)$ and $H(h_\varepsilon)$ now replaced by their counterparts:

$$\tilde{\eta}(h_p; \mu) \equiv \sqrt{\tilde{H}(h_p)}(h + h_p + \mu h_p) / (\mu h_p)$$

and

$$\tilde{H}(h_p; \mu) \equiv \mu h_p (h + h_p) / (h + h_p + \mu h_p).$$

The same argument as in Proposition 3 shows ρ is identified from (A.6), which by the proof of Proposition A1 also holds in this extended model with asymmetric signal precision.

The second step is to identify h and μ . Following the same argument as in the proof of Proposition 3, $\tilde{H}(h_p)$ is identified at $h_p = d_1^{-1}(y_u - y_e)$, and thus the following equation has a known left-hand side.

$$\frac{1}{\tilde{H}(h_p)} = \frac{1}{\mu h_p} + \frac{1}{h + h_p} = \frac{h}{h_p} \left(\frac{1}{\mu h} + \frac{1}{h} \frac{1}{1 + \frac{h}{h_p}} \right). \quad (\text{A.14})$$

Meanwhile, from the definition of $\tilde{M}_p(s_p, h_p)$, we get

$$\text{Var} \left[\tilde{M}_p(s_p, h_p) | h_p \right] = \left(\frac{h_p}{h + h_p} \right)^2 \text{Var} [s_p | h_p] = \frac{1}{h} \left(\frac{1}{1 + \frac{h}{h_p}} \right), \quad (\text{A.15})$$

where the left-hand side is identified at $h_p = d_1^{-1}(y_u - y_e)$ as

$$\text{Var} \left[\tilde{M}_p(s_p, h_p) | h_p = d_1^{-1}(y_u - y_e) \right] = \text{Var} [y_u | y_u - y_e] \equiv v(\Delta y),$$

with $\Delta y \equiv y_u - y_e$. Rearranging (A.15) as

$$\frac{h}{h_p} = \frac{1}{hv(\Delta y)} - 1$$

and plugging it into (A.14) gives

$$I(\Delta y) = \left(\frac{a}{v(\Delta y)} - 1 \right) [b + v(\Delta y)], \quad (\text{A.16})$$

where $I(\Delta y) \equiv 1/\tilde{H}(d_1^{-1}(y_u - y_e))$, and $a \equiv 1/h$, $b \equiv 1/(\mu h)$ are fixed constants. Thus, in this extended model, the equality above must hold as Δy varies continuously over its equilibrium support.

We prove identification of (μ, h) by contradiction. Suppose (μ', h') is observationally equivalent to (μ, h) . Then (A.16) must also hold for (μ', h') . This implies:

$$\frac{a - v(\Delta y)}{a' - v(\Delta y)} = \frac{b' + v(\Delta y)}{b + v(\Delta y)}$$

for all Δy over the support of offer differences in equilibrium. Differentiating both sides above w.r.t. $v(\Delta y)$ and equating the derivatives imply that $(a - a')$ and $(b - b')$ must have the same sign.² But this contradicts the supposition that (A.16) must hold for both (μ, h) and (μ', h') for all Δy . (To see this, note by construction $a/v(\Delta y) - 1$ and $b + v(\Delta y)$ are positive for all (a, b) or (a', b') .) Hence, h and μ are identified.

The next step is to recover the distribution $G_{h_p}(\cdot)$. For each case in the sample, h_p can be recovered from the offer difference Δy using (A.15), where the left-hand side is a directly identifiable conditional variance $v(\Delta y)$. Thus, $G_{h_p}(\cdot)$ is identified nonparametrically.

The final step is to identify m . This follows from the same argument as in the last step in the proof of Proposition 3, only with $M_p(s_p, h_\varepsilon)$ replaced by $\tilde{M}_p(s_p, h_p)$, and $\delta(h_\varepsilon)$ replaced by $\tilde{\delta}(h_p)$. \square

²This argument requires the actual $v(\Delta y)$ in the data-generating process to vary continuously over at least some sections of its equilibrium support. But this follows immediately from continuous variation in h_p in our model.


APPENDIX E. SUBSAMPLE ANALYSIS

In this section, we repeat the estimation and counterfactual analyses after restricting the estimation sample to arbitration cases where both the union and the employer were represented by expert agents. The number of arbitration cases in this subsample is 313. The union is still estimated to be risk-averse, with parameter 0.32. Counterfactual results from the subsample analysis are presented below.

TABLE A3. Conventional Versus Final-Offer Arbitration, 1996-2000

	Conventional, observed	Final-offer, simulated
(a) Mean difference between offers	2.49	1.13
(b) Mean arbitrated wage – offer midpoint	-0.23	0.08
(c) Probability of union win	n/a	0.55
(d) Mean arbitrated wage increase	3.73	3.78

Notes: Column 1 shows average outcomes of the observations in ARB_C . Column 2 Monte Carlo simulates the arbitration model 1000 times conditional on each set of covariates in ARB_C . Offers and wage increases are in units of percentage points.



Notes: Figure displays simulated R^2 values of regression (13) as a function of α , the degree of information transmission. At each value of α , we Monte Carlo simulate 1000 cases per each set of covariates observed in ARB_C and run the regression. For comparison, the dotted, horizontal line marks the R^2 of a regression analogous to (13) run using the observed data from ARB_C . The solid curve and dotted line intersect at $\alpha = 0.27$.

TABLE A4. Efficiency of Awards in CA and FOA

	Conventional $\alpha = 0.37$	Final-offer
$E[-(y - s)^2]$	-0.12	-0.37
$E[- y - s]$	-0.23	-0.42

Notes: The table displays the mean of the efficiency measure across 1000 Monte Carlo simulations conditional on each set of covariates in the ARB_C data set.

TABLE A5. Risk-Averse Versus Risk-Neutral Union in FOA, 1978–1995

	risk neutral	$\rho = 0.32$	$\rho = 1.5$
(a) Mean union offer	8.00	7.74	7.36
(b) Mean employer offer	6.11	6.23	6.30
(c) Probability of union win	0.50	0.56	0.70
(d) Mean arbitrated wage increase	7.05	7.15	7.11
(e) Union's certainty equivalent	7.05	6.75	5.80

Notes: The FOA model is Monte Carlo simulated 1000 times conditional on each set of covariates in the subset of the ARB_F data set where both union and employer were represented by an expert agent. Units are percentage points, excluding probabilities. Employer is risk neutral throughout.

APPENDIX F. MODELS OF SETTLEMENT AND ASSESSING
SELECTION INTO ARBITRATION

Not all collective negotiations of police and fire officer unions in New Jersey are arbitrated. Given that the main data set that we employ for the estimation of our structural model, ARB_F , consists exclusively of cases resolved through arbitration, an interesting question arises of whether selection affects our empirical results. Specifically, if the realization of the parties' signal about the fair wage increase affects the odds that arbitration is required to resolve the dispute, the distribution of signals in our sample could differ from that in the general population of cases.

In this section, we explore this topic through additional theoretical and empirical analysis. We first specify a model of pre-arbitration negotiations that allows the parties to settle their dispute. By settling the case, the parties incur backing-down costs, which reflect a range of factors that affect the desirability of a settlement relative to taking the case to arbitration. Such backing-down costs are orthogonal to the pre-arbitration fair-wage signals perceived by both parties. Our analysis here is quite general, in that we consider alternative versions of the bargaining model—with and without incomplete information between the parties regarding their mutual backing-down costs—which in turn lead to different bargaining solutions. In the case of complete information, our analysis actually accommodates a variety of solutions to the bargaining problem. Our interest is in characterizing the probability of settlement in each bargaining model. Despite the generality of the analysis, we are able to establish, in all versions of the model, a tight connection between the probability that the parties settle the case and the *difference* between their respective certainty equivalents of going to arbitration. Even if the levels of the certainty equivalents change, the settlement probability is not affected if their difference remains the same.

We then turn to our structural model of arbitration to empirically assess how the signal s_p affects the difference between the union's and employer's certainty equivalents of arbitration. Given the estimated model primitives reported in Section 5, we find that changes in the signal received by the parties have essentially no effect on their certainty-equivalent gap. According to our theoretical analysis, it then follows that the conditional probability of going to arbitration is invariant to changes in the signal s_p . Thus, the theoretical and empirical results in the present section show that the estimation results from the main text are fully consistent with a variety of

data-generating processes that produce no systematic selection on signals. We proceed below with the theoretical results in Section F.1 and the empirical evidence in Section F.2.

F.1. A Model of Pre-Arbitration Negotiations. Prior to arbitration, the union and the employer have the opportunity to settle the case. Throughout the present section, we refer to this pre-arbitration interaction as the *negotiation stage*. In the absence of a settlement at the negotiation stage, the case proceeds to the arbitration stage, which consists of the model described in the main text. If the case reaches the arbitration stage, the dispute-resolution process therein eventually gives rise to a wage increase of y . At the negotiation stage, the parties are aware of their own signal, s_p , but they do not know s_a , the signal to be received by the arbitrator if the case proceeds to arbitration. Therefore, from the perspective of the union and the employer at the negotiation stage, y is a random variable.

The union and the employer incur *backing-down* costs c_u and c_e if they settle the dispute prior to arbitration. These costs may vary from one dispute to another. Specifically, for $j \in \{u, e\}$, c_j follows a distribution F_{c_j} . We assume that c_u and c_e are mutually independent from both y and s_p , and that each party knows the realization of its own backing-down costs at the beginning of the negotiation stage. The idea of backing-down costs in negotiations goes back to the classic contributions by Schelling (1956) and Crawford (1982). In our setting of disputes concerning salary increases, Reilly (1963) notes that taking a case all the way to arbitration can often be attractive to the negotiators because it allows them to give their clients the impression of having fought to the end while shifting responsibility to the arbitrator. The backing-down costs, as formulated here, can also include the positive aspects of a settlement, such as the monetary savings associated with the arbitrator and lawyer fees. We interpret backing-down costs flexibly as a term encompassing these various components that affect the desirability of a settlement relative to arbitration.

The negotiation stage payoff structure is as follows: Let $\tilde{y}_u \equiv \frac{-1}{\rho} \log(\mathbb{E}[\exp(-\rho y)])$ be the union's certainty equivalent to obtaining the random wage increase y at the arbitration stage. Similarly, denote by $\tilde{y}_e \equiv \mathbb{E}[y]$ the expected arbitrated wage increase—that is, the negative of the (risk-neutral) employer's certainty equivalent of having y decided at arbitration. If negotiations break down and the parties proceed to arbitration, the expected payoffs, in certainty-equivalent terms, are \tilde{y}_u for the union and $-\tilde{y}_e$ for the employer. If, instead, the parties agree to settle the dispute

for a wage increase of σ , then the payoffs are $\sigma + c_u$ for the union and $-\sigma + c_e$ for the employer.

We next consider two alternative specifications of the negotiation stage game. In one of them, which we refer to as NEG_1 , the realized backing-down costs, c_u and c_e , are common knowledge between the parties. This complete-information environment allows us to be agnostic about the bargaining protocol employed in the pre-arbitration negotiations; we only assume that the solution to the bargaining problem faced by the parties is efficient—that is, disputes only go into arbitration if the overall gains of settling them beforehand are negative. In a second specification, which we refer to as NEG_2 , we consider the case of incomplete information between the parties regarding their backing-down costs. Having incomplete information gives rise to the possibility of inefficient bargaining breakdown, but it also requires us to impose a relatively simple bargaining protocol to keep the analysis tractable.

In both NEG_1 and NEG_2 , we show that any change in the distribution of y that leaves the difference $\tilde{y}_u - \tilde{y}_e$ constant results in no change in the set of realized backing-down costs that are conducive to settlement. Together with the empirical findings presented in Section F.2, these results suggest that the structural estimates in the main text are not heavily affected by non-random selection on signals into our sample of arbitrated cases.

F.1.1. NEG_1 : Backing-Down Costs are Common Knowledge between the Parties. If the backing-down costs are common knowledge between the union and the employer, it is natural to assume that the solution to their pre-arbitration negotiations is efficient. In other words, the parties settle the dispute if and only if their joint gains from settling are positive. Efficiency typically appears in cooperative bargaining solutions that are often adopted by the literature for the analysis of complete information bargaining games. In particular, the Nash bargaining solution (Nash Jr, 1950) has efficiency as one of its axioms. For the purposes of our analysis of the NEG_1 version of the negotiation stage, other than assuming efficiency, we maintain an agnostic view of the exact bargaining protocol adopted by the parties, as well as of the solution to their bargaining problem.

Given our assumption of efficiency, pre-arbitration negotiations break down when the parties' joint gains from settling are negative—that is, if the sum of the settlement payoffs for the union and the employer is smaller than the sum of the parties' certainty

equivalents of going to arbitration. This condition boils down to

$$\sigma + c_u - \sigma + c_e < \tilde{y}_u - \tilde{y}_e,$$

which simplifies to

$$c_u + c_e < \tilde{y}_u - \tilde{y}_e.$$

Under the assumption of independence between the backing-down costs and y , the following proposition is self-evident:

PROPOSITION A4. *Given an initial equilibrium of NEG_1 , the set of backing-down costs leading to arbitration is invariant to a change in the distribution of the potential arbitration awards that holds $\tilde{y}_u - \tilde{y}_e$ constant.*

A direct implication of Proposition A4 is that, upon any change in the distribution of potential arbitration awards that leaves $\tilde{y}_u - \tilde{y}_e$ constant, the equilibrium probability that the dispute reaches the arbitration stage also stays the same.

F.1.2. NEG_2 : Incomplete Information Regarding Backing-Down Costs. We now consider the case in which the backing-down costs, c_u and c_e , are privately known by the parties. Specifically, only the union knows the realization of c_u , and only the employer knows the realization of c_e . We begin by making the following assumption about the distributions of c_u and c_e :

ASSUMPTION A2. *(i) For $j \in \{u, e\}$, F_{c_j} has an associated density function f_{c_j} such that $f_{c_j}(c) > 0$ over its entire support; and (ii) the hazard function associated with the union's cost distribution, $f_{c_u}(c)/[1 - F_{c_u}(c)]$, is strictly increasing in c .*

The bargaining protocol in NEG_2 is take-it-or-leave-it.³ Specifically, the order of play in the negotiation stage is as follows: The union and employer draw their respective costs c_u and c_e . The employer then offers to settle the case for a wage increase σ . If the union rejects the offer, the case proceeds to the arbitration stage.

We solve NEG_2 by backward induction. The union rejects a settlement offer σ if its certainty equivalent of going to arbitration, \tilde{y}_u , is greater than $\sigma + c_u$, its settlement payoff. This condition simplifies to

$$\sigma < \tilde{y}_u - c_u. \tag{A.17}$$

³Though stylized, the take-it-or-leave-it solution is relevant. Perry (1986) shows that in an alternating-offer game with two-sided incomplete information where the cost of bargaining takes the form of a fixed cost per period rather than discounting, the unique sequential equilibrium takes the form of a take-it-or-leave-it offer game.

The employer does not know the union's c_u . Therefore, the employer's problem is

$$\max_{\sigma} F_{c_u}(\tilde{y}_u - \sigma)(-\tilde{y}_e) + [1 - F_{c_u}(\tilde{y}_u - \sigma)](-\sigma + c_e), \quad (\text{A.18})$$

where $F_{c_u}(\tilde{y}_u - \sigma)$ is the probability that the union rejects settlement offer σ . We restrict our attention to interior solutions of the employer's problem—that is, offers that the union accepts with some probability in the interval $(0, 1)$. This restriction is for ease of exposition; but it also comprises the most relevant scenario, given our interest in the selection of cases into arbitration.

The first-order condition associated with (A.18) is

$$\sigma + \frac{1 - F_{c_u}(\tilde{y}_u - \sigma)}{f_{c_u}(\tilde{y}_u - \sigma)} = \tilde{y}_e + c_e, \quad (\text{A.19})$$

where, given A2.(i), the ratio in the left-hand side is defined. The following proposition establishes the property of interest of the equilibrium settlement offer in this specification of the negotiation stage.

PROPOSITION A5. *Consider an initial equilibrium of NEG_2 , as well as a change in the distribution of the potential arbitration award that leaves constant $\tilde{y}_u - \tilde{y}_e$. Then, under Assumption A2, such a change does not affect the set of backing-down costs that leads the union to reject the employer's offer.*

Proof. First apply the change of variable $\tau \equiv \tilde{y}_u - \sigma$ to rewrite (A.19) as

$$\tau - \frac{1 - F_{c_u}(\tau)}{f_{c_u}(\tau)} = \tilde{y}_u - \tilde{y}_e - c_e. \quad (\text{A.20})$$

Hold c_e fixed at an arbitrary value. Then, by assumption, the right-hand side of (A.20) is constant. Meanwhile, Assumption A2.(ii) guarantees that the derivative of the left-hand side of (A.20) with respect to τ is strictly positive. Therefore, τ must be constant as well. Recall that $F_{c_u}(\tau) \equiv F_{c_u}(\tilde{y}_u - \sigma)$ is the probability that the union rejects the employer's offer, resulting in arbitration. Thus, given any c_e , the marginal value of c_u —that is, the value which makes the union indifferent between accepting and rejecting the employer's proposal—is fixed.

□

As in Proposition A4, a direct implication of Proposition A5 is that the equilibrium probability that the case reaches the arbitration stage remains constant given a change in the distribution of the arbitration awards that does not affect the difference $\tilde{y}_u - \tilde{y}_e$.

F.2. Empirical Evidence: The Parties' Signal and the Certainty Equivalent of Arbitration.


We now leverage the structural estimates for the arbitration stage model presented in the main text to assess the relationship between the parties' signal, s_p , and the certainty equivalents of going to arbitration for the union and the employer. With this intent, for each dispute in ARB_F , we compute the parties' certainty equivalents conditional on s_p , based on 10,000 simulated observations.⁴

Figure A6 illustrates the estimated conditional certainty equivalents for one particular dispute in our sample, concerning the wage increase of Atlantic City police officers in the year 1993. As shown in the figure, the relationship between the parties' signal and the union's certainty equivalent is positive—that is, in expectation, the union gets better off by receiving a higher value of s_p . Conversely, the employer expects to pay more as s_p increases. What is notable about the figure is that the union's gains almost exactly compensate the employer's losses, so the gap between the union's certainty equivalent and the negative of the employer's certainty equivalent remains essentially constant as s_p varies.

We now verify whether the pattern concerning the difference in the conditional certainty equivalents between the union and the employer appears more generally throughout our sample. For each dispute in ARB_F , we compute the gap between the parties' certainty equivalents pointwise over a grid of 10,000 values of s_p , and then take the numerical derivative of this gap with respect to s_p . Considering the distribution of the resulting numerical derivatives over all observations in ARB_F and all values of s_p , the 1st, 2.5th, 10th, 90th, 97.5th, and 99th percentiles are -0.05, -0.01, 0.00, 0.00, 0.01 and 0.03, respectively. For reference, the median values of the conditional certainty equivalent levels for the union and the employer across all observations and values of s_p are 7.30 and 7.49, respectively. As these numbers make clear, the derivatives of the certainty-equivalence gap are heavily concentrated in the neighborhood of zero—implying that, in the disputes in our sample, the signal received by the parties does not affect the gap between their certainty equivalents of arbitration in an important manner. Thus, in light of the theoretical results from

⁴To compute the conditional certainty equivalents, we apply the following procedure separately for each of the disputes in ARB_F : first, we draw 10,000 simulated combinations of s , the fair wage increase; and s_p . For each such combination, we compute the parties' final offers and the probability that the union wins arbitration, which suffice for us to obtain the expected payoffs for the union and the employer. We then transform the expected payoffs to obtain the parties' certainty equivalents conditional on s and s_p . Finally, we run kernel regressions of the certainty equivalents on s_p to compute the certainty equivalents conditional on s_p only; in the kernel regressions, we employ the Gaussian kernel and bandwidths given by Silverman's rule of thumb.

FIGURE A6. Certainty equivalents of arbitration conditional on parties' signal



Notes: Figure displays kernel regression estimates of the certainty equivalents of arbitration for the union and the employer on the signal received by the parties, s_p . In the regressions, we employ the gaussian kernel, and the bandwidth is selected by Silverman's rule of thumb. Each regression uses 10,000 simulated observations of the certainty equivalent and s_p , which we compute based on the estimated model primitives for a dispute occurring in 1993 in Atlantic City, NJ. See text and Footnote 4 for details on the simulation.

Section F.1, the empirical findings reported here suggest that the sample of disputes in ARB_F does not suffer from substantial selection on signals of cases into arbitration.