Final-Offer Arbitration

PJ Trainor

1 Introduction

Wage disputes between workers and employers are often settled by arbitration, especially in the case where unions are not allowed to strike. Arbitration comes up not only in salary disputes, but also in situations like medical malpractice cases, and as we will see later, in corporate tax collection. Conventional and final-offer arbitration are the two major forms of arbitration. Both involve two parties, say a union and a firm determining a wage settlement, along with an outside arbitrator. In conventional arbitration, the arbitrator is free to choose any wage settlement. We will focus on understanding final-offer arbitration, in which the two parties make offers to the arbitrator, and then the arbitrator picks one of the offers to be the wage settlement.

We will derive the Nash equilibrium offers that each party makes in the context of final-offer arbitration, but first we need to understand exactly how the arbitrator makes her decision. Firstly, the game starts with the firm and the union making simultaneous offers, which will we denote respectively by w_f and w_u. The arbitrator has a settlement in mind which we will call x, and we assume that she simply chooses whichever offer is closer to x, provided that $w_f < w_u$. [1] If $x < (w_f + w_u)/2$ then she chooses w_f, and if $x > (w_f + w_u)/2$, then she chooses w_u. This is expressed in the figure below, where we consider x to be any value on a real line, with w_f and w_u fixed.

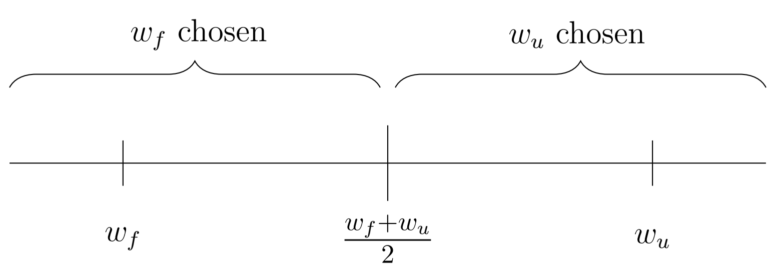

With this machinery in place, we can derive a Nash equilibrium for this game. It will involve basic elements of probability theory like expected value and distribution functions, along with more familiar concepts like optimization problems. Once we have an understanding of the model from a technical standpoint, we will observe how final-offer arbitration is implemented in the context of both settling corporate tax disputes and in settling salary disputes between local governments and police unions.

[1] It can be shown that $w_f < w_u$ is true, but we will not discuss it in this paper.

2 The Model

Considering the terminology we just introduced, the arbiter knows her ideal settlement, x, but the other parties do not [1]. They assume that x is randomly distributed by a cumulative probability distribution $F(x)$, with corresponding probability density function $f(x)$. [2] With these assumptions kept in mind, the parties want to know the probability that their offer is chosen; these probabilities can be expressed as

$$\Pr\{w_f \text{ chosen}\} = \Pr\left\{x < \frac{w_f + w_u}{2}\right\} = F\left(\frac{w_f + w_u}{2}\right).$$

Similarly, we have that

$$\Pr\{w_u \text{ chosen}\} = \Pr\{w_f \text{ not chosen}\} = 1 - F\left(\frac{w_f + w_u}{2}\right).$$

These probabilities tell the firm and the union what chances they have of their offer getting chosen. We can use these probabilities to determine the expected wage settlement, denoted by $E[\text{settlement}]$. Thus we have that

$$\begin{aligned} E[\text{settlement}] &= w_f \Pr\{w_f \text{ chosen}\} + w_u \Pr\{w_u \text{ chosen}\} \\ &= w_f \cdot F\left(\frac{w_f + w_u}{2}\right) + w_u \cdot \left[1 - F\left(\frac{w_f + w_u}{2}\right)\right] \end{aligned}$$

This equation is very important to both the firm and the union; the firm wants to minimize this value, but the union wants to maximize it. If the offer (w_f^*, w_u^*) is to be a Nash equilibrium of this game, then the following must be true:

$$\begin{aligned} w_f^* \text{ solves: } &\min_{w_f} \left\{ w_f \cdot F\left(\frac{w_f + w_u^*}{2}\right) + w_u^* \cdot \left[1 - F\left(\frac{w_f + w_u^*}{2}\right)\right]\right\} \\ w_u^* \text{ solves: } &\min_{w_u} \left\{ w_f^* \cdot F\left(\frac{w_f^* + w_u}{2}\right) + w_u \cdot \left[1 - F\left(\frac{w_f^* + w_u}{2}\right)\right]\right\} \end{aligned}$$

This is something we have seen before—one player solving their own problem assuming the other player has chosen the option that will put them in a Nash equilibrium. This means that the pair (w_f^*, w_u^*) must solve the first order conditions for both of the optimization problems above. These first order conditions are [3]

$$\begin{aligned} (w_u^* - w_f^*) \cdot \frac{1}{2} f\left(\frac{w_f^* + w_u^*}{2}\right) &= F\left(\frac{w_f^* + w_u^*}{2}\right) \\ (w_u^* - w_f^*) \cdot \frac{1}{2} f\left(\frac{w_f^* + w_u^*}{2}\right) &= \left[1 - F\left(\frac{w_f^* + w_u^*}{2}\right)\right] \end{aligned}$$

[2] See the appendix for a brief explanation of distribution functions and how they are treated in this paper.

[3] Note that $\frac{d}{dx}F(x) = f(x)$. This is a property of cumulative distribution functions, and follows from the Fundamental Theorem of Calculus.

Note that the left-hand sides of these equations are equal, which means that if (w_f^*, w_u^*) is a Nash equilibrium, and therefore satisfies the first order conditions, then

$$F\left(\frac{w_f^* + w_u^*}{2}\right) = \left[1 - F\left(\frac{w_f^* + w_u^*}{2}\right)\right] \Longrightarrow F\left(\frac{w_f^* + w_u^*}{2}\right) = \frac{1}{2}$$

This is a very important piece of information, because it tells us that the average of the offers is the median of the arbiter's preferred settlement. This is an impressive result in its own right, but we can go a step further; if we go to the first order condition for the firm, we use the equality above to see that

$$\begin{aligned} F\left(\frac{w_f^* + w_u^*}{2}\right) &= (w_u^* - w_f^*) \cdot \frac{1}{2} f\left(\frac{w_f^* + w_u^*}{2}\right) \\ &= (w_u^* - w_f^*) \cdot F\left(\frac{w_f^* + w_u^*}{2}\right) f\left(\frac{w_f^* + w_u^*}{2}\right) \end{aligned}$$

Then we divide both sides by $F\left(\frac{w_f^*+w_u^*}{2}\right)$ and rearrange terms to get

$$w_u^* - w_f^* = \frac{1}{f\left(\frac{w_f^*+w_u^*}{2}\right)}$$

This means that the difference between the offers must equal the value of the density function at the median of the arbitrator's preferred settlement. That's a complicated statement, but it will make more sense when we consider a concrete example.

Now suppose that the arbiter's preferred settlement is normally distributed with mean μ and variance σ^2. This means that we have the density function

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{\frac{-1}{2\sigma^2}(x-\mu)^2\right\}$$

Since we're looking at a normal distribution, the median is equal to the mean, therefore $(w_f^* + w_u^*)/2 = \mu$, and we see that

$$w_u^* - w_f^* = \frac{1}{f(\mu)} = \sqrt{2\pi\sigma^2},$$

so the Nash equilibrium offers are

$$w_u^* = \mu + \sqrt{\frac{\pi\sigma^2}{2}} \quad \text{and} \quad w_f^* = \mu - \sqrt{\frac{\pi\sigma^2}{2}}$$

When faced with this distribution, the offers are centered around the mean, and as the uncertainty over the arbitrator's preferred settlement increases, so does the gap between the offers. We now have a foundation of our model, and we will see how it operates in the real world.

3 Baseball Arbitration and Tax Disputes

Final-offer arbitration is famously seen in baseball salary arbitration, where a baseball player may choose to bring in a 3rd party to determine their salary, instead of accepting what the team has offered them. In this situation there's the team, the player, and the arbitrator, and once the first two parties have made their offers, the arbitrator picks whichever one she thinks is more appropriate.

The US and Canada face a similar situation when settling corporate tax disputes [3]. When two countries disagree over which one should collect corporate taxes from multinational companies, they can choose enter a final-offer arbitration game, where the winner takes everything and the loser comes out with nothing. The US has beaten Canada over tax disputes three times between 2010 and 2012.

Tax disputes have gone to baseball arbitration with other countries in the past, and the game is attractive for companies because it puts a degree of certainty on its tax bills. Companies can opt for an arbitration agreement if revenue agents have not settled their tax disputes in two years. The way it works is quite simple; a panel is composed of an expert from one country, an expert from the other, and a third person that the first two select. This panel acts as the arbitrator, and revenue agents from both countries submit a tax bill to the panel. Then the panel picks whichever bill it thinks is closest to the correct amount.

We can formulate this in the context of our model. Let r_a be the US's offer, and let r_c be Canada's offer, and let x be the amount that the panel thinks is correct. The two countries are not sure which amount is correct, but they can assume that x is distributed randomly according to a distribution with probability density function $f(x)$ and cumulative density function $F(x)$.[4]

Tax experts from both the US and Canada agree that Canada lost all three corporate tax disputes because it sought too much money in the arbitration. Tax lawyer David Rosenbloom commented that Canada "has developed over the years a habit of taking really extreme and unwarranted positions," and that it's as if Canada is "[almost] unaware arbitration is in the treaty." This means that r_c was much greater than x each time, and Canada should have been playing the game in a different way—they should have lowered their offer in order to get closer to a Nash equilibrium. By Nash equilibrium, we mean that $(r_a + r_c)/2$ is the median of the distribution for x.

Despite being a winner-takes-all game, it should still be played in the same way as the final-offer arbitration game we outlined. While it would be great for Canada if they made a gigantic offer and it was chosen, it might not be realistic, especially if the US is playing differently.

[4] Note that it may not be realistic to assume that the two countries predict that x will follow the same distribution, but for our purposes we will assume they are.

4 New Jersey Police Salaries

Going back to wage disputes, we look at the case of the salaries of New Jersey police officers in the early 1990's [2]. Mayors, local officials, and taxpayers were complaining that police salaries had been rising faster than the rate of inflation, and that they were taking up a large share of the budget in different towns across New Jersey. The state instituted a compulsory-arbitration law over a decade before these complaints started coming up, and it was attributed with the rise in salaries. Much like Canada, it sounds like the persons in charge of these wage disputes should have studied final-offer arbitration before they went to the bargaining table.

To give this example some more context, in 1992, the budget for Glen Rock, NJ was $8.8 million, and salaries made up $3.1 million of that, and 60% of that went to paying police department, whose 20 members compose a fourth of the work force in Glen Rock. In 1988, police salaries only made up 51% of the salary budget. The borough administrator, Robert Freudenrich, was worried about the how necessary items like sewage disposal and health insurance saw decreasing revenues and increased costs, and he was concerned that the residents of the town could not bear tax increases during a recession. There was a budget cap to look out for and the recession lowered property assessments, so the tax base was not what it needed to be.

When the arbitrator was asked to make a decision between the two offers, they were obligated by law to consider the interest of the public, financial consequences, cost of living, and to compare the wages and working conditions of workers in similar areas. Arbitrators were often the target of criticism, but their response was that the cities and town brought weak cases to the bargaining table. Basically, this is another situation where one side did not understand how to play the game, and they were sore losers when they lost. Like in the last example, one party (the government) was offering an extreme proposal, and this resulted in their loss because they made an unreasonable offer. If they wanted to not lose so often, they should have made more reasonable offers, meaning higher ones, or they should have switched to conventional arbitration (alternatively, they could have done some research on final-offer arbitration).

Some states with similar concerns responded with countermeasures, like putting a cap on salary increases or allowing municipalities to reject the final arbitrator's choice, and to do the game over. This second proposal would make final-offer arbitration look like a dynamic game, and it actually can be framed as a game where the parties take turns deciding their offers. We frame it as a static game by assuming that the arbitrator gets both offers simultaneously and that she has an optimal amount in mind. The fact that these measures were proposed and sometimes instituted meant that the police unions understood the game of final-offer arbitration, and the local governments did not, as they hoped the arbitrator would choose their low offers over the unions' more reasonable ones.

5 Conclusion

What we have gone over in this paper is that final-offer arbitration is a relatively simple game with a rather intuitive probabilistic explanation, and it can be understood quite easily without mathematics. Despite the existence of simple explanations, some political bodies like Canadian tax experts or local governments in New Jersey, don't seem to know about them.

While conventional arbitration puts the decision in the hands of a neutral party, and is more likely to produce a settlement that looks more like a "middle ground," final-offer arbitration is much more of a game. The Nash equilibrium of this game occurs when both parties are making reasonable offers, instead of making offers that are too high like Canada did, or ones that are too low like in the case of the New Jersey governments. When one side fails to play the game intelligently, they lose, and in the two cases we looked at, one side lost several times in a row.

The game incorporates uncertainty, as we saw when we looked at the case where x was distributed normally with mean μ and variance σ^2. When uncertainty is higher (higher variance), the parties can make more aggressive offers because more extreme offers are less likely to be in contrast with the arbitrator's preferred settlement. On the other hand, less uncertainty makes it more likely that the arbitrator will prefer settlements closer to μ. Normal distributions are very easy to understand, and it illuminates the implications of the model quite nicely, but the model also works under any probability distribution, and it also assumes that both parties are assigning the same probability distribution to x.

Thus we have developed a setup of this game using reasonable assumptions, and the model is not incredibly complicated, yet it is very useful. We saw two examples that illuminated how this model can be misinterpreted, and that has direct consequences. Perhaps Canada would have gotten some money out of those negotiations, or the New Jersey police salaries would not have been so high if someone had just made a more reasonable offer. Final-offer arbitration is easily implemented in the real world, and it generally does not favor one side, as some critics of it believe, as long as both sides understand how it operates.

6 Bibliography

References

[1] Gibbons, Robert. Game Theory for Applied Economists. Princeton, NJ: Princeton UP, 1992. Print.

[2] Strum, Charles. "Cities Attack Police Raises in New Jersey." The New York Times 14 Feb. 1993, Sunday, Late Edition ed., sec. 1: 41. Print.

[3] Temple-West, Patrick. "International Arbitration for Tax Disputes, Baseball Style." Reuters. Thomson Reuters, 25 Nov. 2012. Web. 21 Mar. 2015.

7 Appendix: Distribution Functions

If we say that a random variable x has probability density function $f(x)$, we mean that $f(x)$ describes the likelihood that the random variable will take on a given value. All probability density functions are nonnegative everywhere and the area under the entire curve is equal to 1. The cumulative distribution function, $F(x)$, takes in a value x^*, and it tells us the probability that the random variable will take on a value less than or equal to x^*. More specifically,

$$F(x^*) = \int_{-\infty}^{x^*} f(x)dx.$$

This is also presented in the picture below. Furthermore, we can use the fundamental theorem of calculus to see that

$$\frac{d}{dx}F(x) = f(x).$$

This is all we need to know about distribution functions in order to understand final-offer arbitration.