

NBER WORKING PAPER SERIES

LEARNING IN EQUILIBRIUM MODELS OF ARBITRATION

Robert S. Gibbons

Working Paper No. 2547

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 1988

I am happy to acknowledge several forms of assistance and support: I have benefitted greatly from extensive discussions about arbitration with David Bloom, Christopher Cavanagh, Henry Farber, and Stanley Jacks. Many others have made helpful comments on this paper: David Card, Vincent Crawford, Paul Klemperer, Gregory Leonard, Meg Meyer, Keven M. Murphy, Joel Sobel, and seminar audiences at Chicago, Harvard, MIT, and Princeton. Finally, the NSF provided financial support through grant IST 86-09691. The research reported here is part of the NBER's research program in Labor Studies. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

Learning in Equilibrium Models of Arbitration

ABSTRACT

This paper analyzes strategic communication in equilibrium models of conventional and final-offer interest arbitration. Both models emphasize the role of learning by the arbitrator from the parties' offers about the state of the employment relationship, which is known to the parties but not to the arbitrator. In both models, the arbitrator's equilibrium behavior is identical to the reduced-form decision rule typically assumed in the empirical literature. The paper thereby provides a structural interpretation for the existing empirical work.

The paper also represents progress towards a complete theory of arbitration because it satisfies three conditions that will be required of any such theory. First, the models' predictions match the existing empirical evidence. Second, the models describe equilibrium behavior. And third, the models are built on a common set of assumptions about preferences, information, and commitment. The paper therefore not only provides an equilibrium foundation for the intuition that the arbitrator might learn from the parties' offers, but also uses the idea of learning to develop a unified analytical treatment of the two major forms of interest arbitration.

Robert S. Gibbons
Department of Economics
E52-251B
MIT
77 Massachusetts Avenue
Cambridge, MA 02139

1. Introduction

This paper analyzes strategic communication in equilibrium models of conventional and final-offer interest arbitration. Both models emphasize the role of learning by the arbitrator from the parties' offers about the state of the employment relationship, which is known to the parties but not to the arbitrator. In both models, the arbitrator's equilibrium behavior is identical to the reduced-form decision rule typically assumed in the empirical literature (namely, a weighted average of the parties' offers and a settlement motivated by the publicly observable facts of the case). The paper thereby provides a structural interpretation for the existing empirical work.

The paper also represents progress towards a complete theory of arbitration because it satisfies three conditions that will be required of any such theory. First, the models' predictions match the existing empirical evidence. Second, the models describe equilibrium behavior. And third, the models are built on a common set of assumptions about preferences, information, and commitment. (This last condition is crucial for accurate welfare comparisons of conventional and final-offer arbitration.) The paper therefore not only provides an equilibrium foundation for the intuition that the arbitrator might learn from the parties' offers, but also uses the idea of learning to develop a unified analytical treatment of the two major forms of interest arbitration.

Understanding interest arbitration is important for several reasons. First, considerable resources are allocated through interest arbitration of wage disputes, both in the public sector (especially where strikes are prohibited) as well as in private-sector industries such as professional baseball. Second, outside the context of wage disputes, the clogged judicial system has created a growing interest in alternative dispute-resolution procedures, especially pre-trial arbitration. Third, interest arbitration offers a rare natural

experiment involving strategic behavior and significant stakes: the rules of arbitration constitute a well defined game, and the players' choices often are observed without error.

A fourth reason to understand arbitration is that it is an institution that allows parties to write an incomplete contract *ex-ante* while limiting opportunistic behavior *ex-post*. Indeed, Williamson (1975) argues that the possibility of arbitration by informed insiders is one of the central virtues of internal over market organization. Even if some state-contingent actions cannot be specified in a contract *ex-ante*, the contract can include the provision that *ex-post* disputes will be resolved through arbitration. Grossman and Hart (1986) analyze two other responses to incomplete contracts: integration (i.e., assigning residual rights of control) and *ex-post* bargaining. Arbitration is a third option.

Formally, this paper considers interest arbitration, as opposed to grievance arbitration, but the thrust of the paper seems applicable to grievance arbitration as well. In interest arbitration the arbitrator settles wage (or other) disputes that arise when a new contract is negotiated; in grievance arbitration, the arbitrator settles disputes that arise during the life of an existing contract. The two major forms of interest arbitration are conventional arbitration and final-offer arbitration. In final-offer arbitration, the parties simultaneously submit wage offers, and then the arbitrator imposes one of the two offers as the settlement; in conventional arbitration, the arbitrator's choice of a settlement is unconstrained.

Although this paper contains no data, it is closely related to the empirical literature on interest arbitration. As noted briefly above, the arbitrator's equilibrium behavior in this paper is identical to the reduced-form decision rule typically assumed in the empirical literature. In both the conventional and the final-offer arbitration models developed below, the arbitrator's preferred settlement, y_a , is a weighted average of the average of the parties'

offers, $\bar{y} = (y_e + y_o)/2$, and a settlement motivated by the publicly observable facts of the case, $y^*(f)$:

$$y_a = \alpha \bar{y} + (1-\alpha) y^*(f).$$

A decision rule of this form appears in empirical studies of conventional arbitration as equation (12) in Ashenfelter and Bloom (1984), equation (3) in Farber and Bazerman (1986), and equation (1b) in Bloom (1986); it also appears in empirical studies of final-offer arbitration as equation (7) in Ashenfelter and Bloom and equation (6) in Farber and Bazerman. One of the purposes of this paper is to provide a structural interpretation for the estimated parameter α .

The paper also is an attempt to move towards a complete theory of arbitration. A first requirement for such a theory is that it match the empirical evidence. Perhaps the strongest piece of evidence is that the parties' offers matter in the arbitrator's decision-making (i.e., α is positive; see Bloom and Farber-Bazerman). Indeed, Stevens' (1966) proposal of final-offer arbitration was motivated by the observation that the offers seem to matter too much in conventional arbitration: the arbitrator sometimes seems simply to split the difference between the parties' offers, ignoring the publicly observable facts of the case. This paper uses the arbitrator's need to learn from the parties' offers to explain the observed importance of the offers in the arbitrator's decision-making. The equilibrium behavior that emerges also is consistent with the Farber-Bazerman finding that the importance of the offers in the arbitrator's decision-making increases when the offers are closer together, and with the Ashenfelter-Bloom and Farber-Bazerman finding that the

arbitrator's notion of an ideal settlement is identical in final-offer and conventional arbitration.

A second requirement for a complete theory of arbitration is that it describe equilibrium behavior. The following discussion explains why an equilibrium analysis of arbitration is a crucial complement to empirical work on the arbitrator's decision rule. Consider Stevens' observation that in conventional arbitration the arbitrator sometimes seems simply to split the difference between the offers. If the arbitrator were committed to splitting the difference, then the parties would submit the most extreme offers possible. Suppose instead that the arbitrator is committed to splitting the difference between reasonable offers but to ignoring (or even punishing) unreasonable ones. Such a decision rule could cause the parties always to submit reasonable offers, and so lead to the appearance of mechanical compromising by the arbitrator, but the observed behavior would be only half the story: the data would be silent about how the arbitrator would have behaved had the parties not submitted reasonable offers. Out-of-equilibrium responses by the arbitrator are never observed, and so can never be estimated, but play a crucial role in determining what is observed. Only equilibrium models can elucidate the out-of-equilibrium behavior that empirical analyses are forced to ignore.

A third requirement for a complete theory of arbitration is that models of different forms of arbitration be built on a common set of assumptions about preferences, information, and commitment. This is a prerequisite for accurate comparisons of conventional and final-offer arbitration for two reasons, the first empirical and the second theoretical. First, if each form is adopted in environments in which it performs well (and if the data cannot reveal the differences in these environments), then comparisons of performance data from real arbitrations are likely to suffer from selection bias. And

second, because any form of arbitration could in principle be adopted in any given environment, a successful model of a given form of arbitration must be built on assumptions that permit other forms of arbitration to be analyzed within the same model. Consider Farber's (1980) model of final-offer arbitration, for example. Farber's model is the leading contribution to the existing theoretical literature on that form of arbitration, and is the basis for all of the empirical work on the subject. Its important liability, however, is that the analogous model of conventional arbitration implies that the parties' offers are irrelevant in the arbitrator's decision-making, an implication that is strongly contradicted by the empirical evidence.

This paper is able to meet these three requirements because it emphasizes the arbitrator's attempt to learn from the parties' offers. The paper introduces the possibility of such learning by assuming that the arbitrator's preferred settlement is determined not only by the publicly observable facts of the case but also by the state of the employment relationship, which is known to the parties but not to the arbitrator. The most important distinction between the facts of the case and the state of the employment relationship is simply that the arbitrator observes the former but not the latter. Examples of publicly observable facts are the inflation rate, recent wage settlements in related industries, a coarse indication of the firm's financial health, and so on. The state of the employment relationship, in contrast, is a collection of privately observable details of the production environment and the labor market that the employer and union inhabit. Because of their long and close relationship, the employer and the union understand the state well. The arbitrator, in contrast, has less experience with these details but conducts an independent investigation to gather information.

Two important assumptions maintained throughout the paper deserve brief mention. First, the arbitrator is assumed to be unable to commit to a decision rule before the parties submit their offers. This assumption seems to represent observed practice accurately, but the complementary case also deserves attention. Second, the parties are assumed to share a common perception of the state of the employment relationship. While it seems realistic that the parties should both be considerably better informed about the state than the arbitrator is, it also seems likely that the parties themselves will have different beliefs about the state. Detailed discussions of both of these assumptions are provided in Section 7.

The remainder of the paper is organized as follows. Section 2 describes Farber's equilibrium model of final-offer arbitration. As will become clear, this model plays an important role in the development of the learning models of both final-offer and conventional arbitration analyzed here. Section 3 specifies the economic environment (i.e., makes assumptions about the players' preferences and information structures) that underlies this paper's models of both final-offer and conventional arbitration. Sections 4 and 5 analyze separating equilibria in the final-offer and conventional arbitration games, respectively. In both cases, the parties' offers perfectly reveal their private information, and the arbitrator uses this information to determine an ideal settlement. In final-offer arbitration the arbitrator chooses the offer that is closer to this ideal as the settlement; in conventional arbitration the arbitrator simply imposes the ideal itself. The model of conventional arbitration in Section 5 is of independent interest as a contribution to the theory of cheap-talk communication games. Section 6 summarizes the relationship between these learning models and the empirical literature. Finally, Section 7 discusses a program of extensions of the models developed here.

2. Final-Offer Arbitration without Learning

This section presents the key points in Farber's model of final-offer arbitration. The analysis of learning in equilibrium models of final-offer and conventional arbitration that follows in Sections 4 and 5 builds on the work described here.

Farber characterizes an arbitrator by a number, z , that describes the arbitrator's most preferred settlement in the case in question. To facilitate comparison of this model with those that follow, suppose that the arbitrator's utility if the ideal settlement is z and the actual settlement is y is given by the utility function $v_a(y,z) = -(y-z)^2$. In final-offer arbitration, the arbitrator is constrained to choose one of the parties' offers as a settlement. Assuming that the employer's offer, y_e , is below the union's offer, y_u , the arbitrator chooses the employer's offer if and only if $z < \bar{y}$, where $\bar{y} = (y_e + y_u)/2$ is the average of the parties' offers.

The parties are assumed to be uncertain about the value of z . Let this uncertainty be described by the probability distribution $F(z)$. Assume that $F(z)$ has a continuous and strictly positive density $f(z)$ with connected support. Assume also that the parties have strictly opposed preferences and are risk-neutral: the employer simply seeks to minimize the arbitrator's expected settlement, while the union seeks to maximize it.

The timing of the arbitration game is as follows. First, the parties simultaneously submit their offers to the arbitrator. Then the arbitrator chooses the offer that maximizes the utility function $v_a(y,z)$. The parties' Nash equilibrium offers therefore must solve

$$(1) \quad \max_{y_e} y_e F(\bar{y}) + y_u [1-F(\bar{y})], \quad \text{and} \\ \max_{y_u} y_u [1-F(\bar{y})] + y_e F(\bar{y})$$

$$(2) \quad \min_{y_e} \quad y_e F(\bar{y}) + y_u [1 - F(\bar{y})].$$

The first-order conditions for these optimization problems are

$$(3) \quad 1 - F(\bar{y}) = \frac{1}{2} (y_u - y_e) f(\bar{y}) \quad \text{and}$$

$$(4) \quad F(\bar{y}) = \frac{1}{2} (y_u - y_e) f(\bar{y}),$$

which imply that $F(\bar{y}) = 1/2$ and $y_u - y_e = f(\bar{y})^{-1}$. (Note that the assumption that $y_u > y_e$ clearly cannot be false in equilibrium.) These first-order conditions express a simple trade-off: in equilibrium, each party considers making a more aggressive offer but must balance the gain from having such an offer accepted against the reduced probability that it will be accepted.

Much of what follows assumes that z is Normally distributed with mean M and precision H (i.e., variance $1/H$). In this case, (3) and (4) imply that the equilibrium offers are

$$(5) \quad y_u = M + \sqrt{\frac{\pi}{2H}} \quad \text{and}$$

$$(6) \quad y_e = M - \sqrt{\frac{\pi}{2H}}.$$

(Checking for the global optimality of these and subsequent offers is tedious and so is relegated to Appendix 1.) Note that the equilibrium offers are centered about the mean of the parties' belief about the arbitrator's ideal settlement, and that the distance between the equilibrium offers decreases as this belief becomes more precise.

3. An Economic Environment

This section describes the preferences and information structure that underlie both the final-offer and the conventional arbitration models analyzed in Sections 4 and 5. The novel feature introduced here is an unknown called the state of the employment relationship (or simply the state) and denoted by s . The state of the employment relationship represents the details of the production environment and the labor market that the employer and the union inhabit. Because of their long and close relationship, the employer and the union understand the state well. The arbitrator, in contrast, has less experience with these details but conducts an independent investigation to gather information.

The central theme of the models analyzed below is as follows. If the arbitrator knew the state precisely then the ideal settlement would be clear. Because the arbitrator is comparatively poorly informed, however, any information about the state that can be extracted from the parties' offers may be useful. Such inferences must of course account for the natural incentive of each party to mislead the arbitrator.

3a. Preferences

Assume that the arbitrator's ideal settlement depends on the state of the employment relationship: $z = z(s)$. This dependence could arise because the arbitrator wishes to be fair,

or to be rehired. For simplicity, let $z(s)=s$ and let the arbitrator's (von Neumann-Morgenstern) utility function be $v_a(y,s) = -(y-s)^2$.

Unlike the arbitrator, the parties are assumed to be risk-neutral and to have preferences that are independent of the state: $v_e(y,s) = -y$ and $v_u(y,s)=y$. As in the previous section, the employer simply seeks to minimize the arbitrator's expected settlement, while the union seeks to maximize it. Thus, the parties are well informed about what a fair settlement would be, but this information does not affect their preferences.

These assumptions about preferences keep the analysis simple. It may seem odd, however, that the state s appears in the arbitrator's utility function but not in the parties' utility functions. Two justifications for this can be given. First, the intuition developed here seems unlikely to disappear when the parties' preferences do depend on the state. And second, the assumptions could be correct as they stand. Let the value of the enterprise be $V(s)$ when the state is s , let the settlement y be the wage bill, and let employment be fixed. Then the union's utility function is $v_u(y,s)=y$, the employer's is $v_e(y,s) = V(s) - y$ (which changes nothing in the analysis), and the arbitrator's ideal settlement $z(s)=s$ reflects a judgment about the fair division of $V(s)$.

3b. Information Structure

Let s be Normally distributed with mean m and precision h . Also, let the parties' knowledge of the state be summarized by the noisy signal

$$(7) \quad s_p = s + \epsilon_p,$$

where ϵ_p is Normally distributed with zero mean and precision h_p . Notice that the parties share a common perception of the state of the employment relationship. This is important in what follows, and is discussed in Section 7. Finally, let the results of the arbitrator's private investigation be represented by the noisy signal

$$(8) \quad s_a = s + \epsilon_a,$$

where ϵ_a is Normally distributed with zero mean and precision h_a . The three random variables s , ϵ_p , and ϵ_a are assumed to be independent of each other. In keeping with the idea that the parties are better informed about the state than is the arbitrator, one could assume that $h_p > h_a$, but this is not necessary for what follows.

To summarize, the information structure is as follows: the parties both observe s_p , the arbitrator observes s_a , no one observes s , and everyone observes m , h , h_p , and h_a . The parameters m and h that specify the prior distribution of s can be interpreted as the publicly observable facts of the case.

Because of the independence of s , ϵ_p , and ϵ_a , the simple formulae of the Normal learning model (DeGroot, 1970, Chapter 9) characterize the Bayesian updating that occurs in the model. Some of the results of this updating are recorded here for future use. The conditional distribution of s given s_i (where $i \in \{p, a\}$) is Normal with mean $M_i(s_i)$ and precision H_i , where

$$(9) \quad M_i(s_i) = \frac{hm + h_i s_i}{h + h_i} \quad \text{and}$$

$$(10) \quad H_i = h + h_i.$$

Similarly, the conditional distribution of s given s_p and s_a is Normal with mean $M_{pa}(s_p, s_a)$ and precision H_{pa} , where

$$(11) \quad M_{pa}(s_p, s_a) = \frac{hm + h_p s_p + h_a s_a}{h + h_p + h_a} \quad \text{and}$$

$$(12) \quad H_{pa} = h + h_p + h_a.$$

These definitions provide a convenient shorthand for describing the beliefs and decision-making in the models below.

4. Learning in Final-Offer Arbitration

This section analyzes a model of final-offer arbitration in the economic environment specified in Section 3. The analysis focuses on a separating equilibrium: The arbitrator perfectly infers the parties' private information (s_p) from their offers and then uses this information and the signal (s_a) to compute a posterior belief about the state (s) and to choose the expected-utility-maximizing offer as the settlement. The parties understand that the arbitrator will draw an inference from the offers and so consider the payoff from misleading the arbitrator when choosing their offers, but find it optimal not to submit misleading offers.

The timing of the game is as follows. First, the parties observe s_p and the arbitrator observes s_a . Second, the parties simultaneously submit offers, y_e and y_u . Third, the arbitrator chooses the offer that maximizes $E_s[v_a(y, s) \mid s_a, y_e, y_u]$, where the conditioning on

y_e and y_u denotes the information about s_p that the arbitrator extracts from the parties' offers. Note that the arbitrator is not allowed to commit to a decision rule in advance; rather, the arbitrator's choice must satisfy sequential rationality at the third stage of the game.

Before describing the separating equilibrium in detail, it helps set the tone for the subsequent analysis to show that some learning must occur in equilibrium. That is, it cannot be equilibrium behavior for the arbitrator to learn nothing from the parties' offers. To see this, suppose the arbitrator thought that no information could be extracted from the offers. Then the arbitrator's ideal settlement, $y_a(s_a)$, would be $M_a(s_a)$ from (9), or

$$(13) \quad y_a(s_a) = \frac{hm + h_a s_a}{h + h_a}.$$

Given s_p , the parties' belief about $y_a(s_a)$ is Normally distributed with mean m' and precision h' , where

$$(14) \quad m' = \frac{hm + h_a M_p(s_p)}{h + h_a} \quad \text{and}$$

$$(15) \quad h' = \frac{(h + h_p)(h + h_a)^2}{h_a(h + h_p + h_a)}.$$

(Note that the dependence of m' on s_p has been suppressed in the notation.) The logic behind Farber's equilibrium now applies directly: if the arbitrator learns nothing from the parties' offers then s_p is relevant only because it changes the parties' beliefs about the arbitrator's ideal settlement. The equilibrium offers are therefore analogous to (5) and (6):

$$(16) \quad y_u(s_p) = m' + \sqrt{\frac{\pi}{2h'}} \quad \text{and}$$

$$(17) \quad y_e(s_p) = m' - \sqrt{\frac{\pi}{2h'}}.$$

Contrary to the hypothesis, the arbitrator can learn something from these offers. The average of the offers is m' , and (9) and (14) determine s_p from m' . Let $s_p(m')$ denote the (point) estimate of s_p computed in this way. Sequential rationality then requires that the arbitrator's ideal settlement be $M_{pa}(s_p(m'), s_a)$ from (11), or

$$(18) \quad y_a(y_e, y_u, s_a) = \frac{hm + h_p s_p(m') + h_a s_a}{h + h_p + h_a}.$$

The fact that (18) differs from (13) establishes the following result.

Proposition 1. There does not exist a perfect Bayesian equilibrium in which the arbitrator learns nothing from the parties' offers.

The crucial difference between (13) and (18), of course, is that in the latter the parties' offers help the arbitrator determine the ideal settlement. In an equilibrium featuring this kind of learning, each party's offer must balance not only the probability of and the gain from submitting the offer chosen as the settlement, as in Farber's model, but also the effect the offer has on the arbitrator's inference about the ideal settlement. The rest of this section establishes that a separating perfect Bayesian equilibrium exists in which the parties' offers perfectly reveal s_p to the arbitrator.

Suppose that the arbitrator believes that \bar{y} , the average of the parties' offers, perfectly reveals s_p . That is, for any pair of offers y_e and y_u the arbitrator computes the point estimate $s_p = s_p(\bar{y})$. (Other separating equilibria may exist in which some other function of y_e and y_u reveals s_p to the arbitrator.) The arbitrator's ideal settlement is then

$$(19) \quad y_a(y_e, y_u, s_a) = \frac{hm + h_p s_p(\bar{y}) + h_a s_a}{h + h_p + h_a},$$

and (assuming $y_e < y_u$) the arbitrator chooses y_e as the settlement if and only if the ideal settlement, $y_a(y_e, y_u, s_a)$, is less than the average of the offers, \bar{y} .

It is simple to show that if the equilibrium value of \bar{y} is to reveal s_p in this way then the arbitrator's inference rule must be

$$(20) \quad s_p(\bar{y}) = \frac{(h + h_p)\bar{y} - hm}{h_p}.$$

To see this, use (19) and (11) to express the event that the arbitrator chooses y_e as $s_a < S(\bar{y})$, where

$$(21) \quad S(\bar{y}) = \frac{h_a \bar{y} + h(\bar{y} - m) + h_p(\bar{y} - s_p(\bar{y}))}{h_a}.$$

A derivation analogous to that leading to (3) and (4) then yields the first-order conditions

$$(22) \quad 1 - F[S(\bar{y})] = (y_u - y_e) f[S(\bar{y})] S'(\bar{y}) \quad \text{and}$$

$$(23) \quad F[S(\bar{y})] = (y_u - y_e) f[S(\bar{y})]S'(\bar{y}),$$

where F is now the distribution of s_a conditional on s_p , which is Normal with mean $M_p(s_p)$ given by (9) and precision

$$(24) \quad H' = \frac{h_a(h + h_p)}{h + h_p + h_a}.$$

As in (5) and (6), these first-order conditions imply that $S(\bar{y}) = M_p(s_p)$. Substituting $s_p(\bar{y})$ for s_p in $M_p(s_p)$ and solving for $s_p(\bar{y})$ yields (20), the desired result.

The inference rule $s_p(\bar{y})$ in (20) implies that (21) simplifies to $S(\bar{y}) = \bar{y}$: the arbitrator's decision rule is to choose y_e if and only if $s_a < \bar{y}$. Note that this is directly analogous to the decision rule used in Farber's model (where there is no learning). Not surprisingly, therefore, the equilibrium strategies for the parties that solve the first-order conditions (22) and (23) are directly analogous to (5) and (6):

$$(25) \quad y_u(s_p) = M_p(s_p) + \sqrt{\frac{\pi}{2H'}} \quad \text{and}$$

$$(26) \quad y_e(s_p) = M_p(s_p) - \sqrt{\frac{\pi}{2H'}}.$$

The average of the offers is $M_p(s_p)$, which yields a more convenient expression of the inference rule (20): in equilibrium, the arbitrator infers s_p from the average of the parties' offers, uses it in (19) to compute the ideal settlement, and then chooses the offer closer to this ideal.

The summary of this argument is the main result of this section.

Proposition 2. The parties' offer strategies given in (25) and (26) and the arbitrator's decision strategy based on the ideal settlement (19) and the inference rule (20) constitute a separating perfect Bayesian equilibrium of the final-offer arbitration game. In this equilibrium, the arbitrator's ideal settlement can be written

$$y_a = \alpha \bar{y} + (1-\alpha) s_a, \quad \text{where}$$

$$\bar{y} = \frac{hm + h_p s_p}{h + h_p} \quad \text{and}$$

$$\alpha = \frac{h + h_p}{h + h_p + h_a}.$$

The key feature of the arbitrator's equilibrium strategy is its treatment of unexpected offers from the parties. In equilibrium, the parties' offers differ by $\sqrt{2\pi/H'}$, and the arbitrator perfectly infers the parties' private information from the average of the offers. Any pair of offers that differs by this prescribed amount is on the equilibrium path; a pair that differs by any other amount is off the equilibrium path. The arbitrator's inference rule ignores this distinction, however: (20) produces the point estimate of s_p both on and off the equilibrium path. Thus, any change in either party's offer changes the arbitrator's estimate of s_p at the rate $(h + h_p) / 2h_p$. The reason the arbitrator is not misled in equilibrium is that the parties must balance this opportunity to influence the arbitrator's belief against the two considerations that determined the equilibrium offers in Farber's model, the gain from having a more aggressive offer accepted and the reduced probability that it will be accepted.

There are two cases in which learning plays no role in the private-information game analyzed in this section: (1) the arbitrator has nothing to learn (i.e., h_a approaches infinity);

and (2) the parties have nothing to communicate (i.e., h_p approaches zero). One appealing feature of the separating equilibrium analyzed here is that it approaches the equilibrium in Farber's complete-information game as the learning becomes unimportant in the private-information game. As h_a approaches infinity, for instance, the arbitrator's uncertainty about s disappears, and s_p becomes relevant only as a signal to the parties about the arbitrator's ideal settlement. In this case, the precision of the parties' belief about the arbitrator's ideal settlement (i.e., H in (5) and (6) in Farber's model) is $h + h_p$, which is indeed the limit of H' in (25) and (26) as h_a approaches infinity. Similarly, as h_p approaches zero, the parties' private information does not improve their prior information about the arbitrator's ideal settlement, so H in (5) and (6) is $hh_a/(h + h_a)$, which is indeed the limit of H' as h_p approaches zero.

One other limit also bears consideration. As h_p approaches infinity, the arbitrator perfectly infers not only s_p but also s from the average of the parties' offers. Given perfect information about s , the arbitrator's private signal s_a becomes irrelevant in determining the arbitrator's ideal settlement, and this eliminates the parties' uncertainty about the arbitrator's ideal settlement. Recall that as the analogous uncertainty disappears in Farber's complete-information model, the parties' equilibrium offers converge to the mean of the distribution of the arbitrator's ideal settlement. (Formally, as H approaches infinity, the equilibrium offers given in (5) and (6) converge to M .) Interestingly, as h_p approaches infinity, H' does not approach infinity (rather, it approaches h_a), so the equilibrium offers given by (25) and (26) in the private-information model do not converge. This happens because (as noted above) (20) and (21) imply that the arbitrator chooses y_e if and only if $s_a < \bar{y}$, independent of the value of h_p . Thus, s_a serves as a tie-breaker in the arbitrator's decision rule even though it is irrelevant to the arbitrator's inference process.

5. Learning in Conventional Arbitration

This section analyzes a model of conventional arbitration in the economic environment specified in Section 3. The analysis focuses on a continuum of payoff-equivalent separating equilibria. In each of these equilibria, the arbitrator perfectly infers the parties' private information (s_p) from their offers and then uses this information and the signal (s_a) to compute a posterior belief about the state (s) and to impose whatever settlement maximizes expected utility given this belief. The parties understand that the arbitrator will draw an inference from the offers and so consider the payoff from misleading the arbitrator when choosing their offers, but find it optimal not to submit misleading offers.

The timing of the game is as follows. First, the parties observe s_p and the arbitrator observes s_a . Second, the parties simultaneously submit offers, y_e and y_u . Third, the arbitrator imposes whatever settlement maximizes $E_s[v_a(y,s) | s_a, y_e, y_u]$. As in the previous section, the arbitrator is not allowed to commit to a decision rule in advance. It will become clear below, however, that in this model of conventional arbitration the arbitrator would not value the opportunity to commit to a decision rule: each of the continuum of sequentially rational decision rules associated with the continuum of separating equilibria described below would be an optimal rule if the arbitrator could commit.

The crucial difference between this model of conventional arbitration and the previous section's model of final-offer arbitration is that here arbitration is modeled as a cheap-talk game: the parties' offers are not directly relevant to any player's payoff; they matter only indirectly through the information they convey.¹ In this case, it is natural in the model (as well as faithful to the institutional details of arbitration) to envision party i 's offer

¹See Crawford and Sobel (1982) for the seminal work on cheap-talk games.

as a long speech that ends "and that is why we believe that an appropriate settlement is y_1 ." Final-offer arbitration is not a cheap-talk game because one of the parties' offers must be chosen as the settlement, and so directly affects payoffs. Similarly, conventional arbitration would not be a cheap-talk game if the offers in the arbitration game were constrained to be the parties' last offers in a prior bargaining game, because such offers could have been accepted and would then have directly affected payoffs.²

All cheap-talk games share two distinctive features. The first is that in any cheap-talk game there exists a trivial equilibrium in which no communication occurs. In the conventional arbitration model, for instance, the arbitrator expects there to be no information in either party's offer and so pays no attention to either offer when choosing a settlement, and each party is content to choose an offer at random because no offer has any affect on the arbitrator's settlement. Thus, there is no analog of Proposition 1 here: in this model of conventional arbitration there does exist an equilibrium in which the arbitrator learns nothing from the parties' offers. Note well, however, that in this equilibrium the arbitrator's ideal (and imposed) settlement is independent of the parties' offers. Since such independence is not observed in the data, this no-communication equilibrium is suspect.

The second distinctive feature of all cheap-talk games is that, given a fixed cheap-talk game, any perfect Bayesian equilibrium (separating or otherwise) is associated with a class of payoff-equivalent perfect Bayesian equilibria that differ only by a translation: since messages are costless, the information communicated in any given equilibrium also can be communicated in other payoff-equivalent equilibria that differ only in the language used to

²Thus, in this paper, the parties' offers in arbitration need not be identical to their last offers in a prior bargaining game. This flexibility arises if (as is required in some contracts) the bargaining record has been sealed and is unavailable to the arbitrator. Also, it will become clear below that the incentive for the parties to make concessions in arbitration can be quite strong, perhaps stronger than the analogous incentive in bargaining. Therefore, even if the bargaining record has not been sealed, the last offers from the bargaining game may be sufficiently far apart that they have no bearing on the parties' choices of offers in arbitration.

communicate. This accounts for the great multiplicity of payoff-equivalent separating equilibria mentioned above.

This section's model of conventional arbitration is a cheap-talk game with a third distinctive feature: neither party alone can achieve any communication with the arbitrator, because each party has preferences over settlements that are independent of the state of the employment relationship (the employer simply seeks to minimize the arbitrator's expected settlement, while the union seeks to maximize it). This section demonstrates, however, that together the parties can achieve perfect communication: there exist separating perfect Bayesian equilibria in which the parties' offers perfectly reveal s_p to the arbitrator.

One scheme for extracting a single piece of information from two informed agents is to ask each to make a claim about the information and to shoot them both if their claims disagree.³ Such schemes are not feasible in the arbitration model analyzed here, however, for two reasons. First, the arbitrator's power is limited to imposing a transfer between two risk-neutral parties, so outcomes that are Pareto-inefficient for the two parties are impossible. And second, the arbitrator's behavior must be sequentially rational: given any pair of offers (y_e, y_u) and any signal (s_a) , the arbitrator must formulate a belief about the parties' private information, $\mu(s_p | s_a, y_e, y_u)$, and act optimally given this belief.

In spite of the facts that Pareto-inefficient outcomes are impossible and that the arbitrator must be sequentially rational, the arbitrator can extract s_p from the parties. The crucial feature of the model that drives this result is that the parties' information (s_p) is correlated with the signal from the arbitrator's investigation (s_a).

To understand this result, it is helpful to consider the intermediate case in which Pareto-inefficient outcomes are impossible but sequential rationality is not required. In this

³See Kalai and Rosenthal (1979) for a more careful (if less dramatic) statement of this result.

case, the following scheme induces truth-telling: ask each party for a claim about s_p ; if they disagree then compare each party's claim about s_p to the signal s_a and enforce a large transfer to the party who seems to have made the more reasonable claim from the party who seems to have made the less reasonable claim. Given any correlation between s_p and s_a , there exists a sufficiently large transfer such that truth-telling is the unique equilibrium of this scheme.

It is simple to extend the spirit of this scheme to the model of conventional arbitration studied here, in which the arbitrator's behavior must be sequentially rational. The argument relies on the fact that in a perfect Bayesian equilibrium of the conventional arbitration game, the arbitrator's belief off the equilibrium path is unrestricted. This freedom ensures that there exists a belief off the equilibrium path that makes it sequentially rational for the arbitrator to enforce the required transfer if the parties disagree.

There is a close relationship between the incentives created in a conventional arbitration game in which the arbitrator uses the scheme described above and the incentives created in a final-offer arbitration game: the size of the transfer and the conditions under which it is awarded to one party or the other in the conventional arbitration game are analogous to the gap between the offers and the conditions under which the arbitrator chooses one party's offer or the other's in the final-offer arbitration game. In fact, the proof that there exists a continuum of separating perfect Bayesian equilibria in the conventional arbitration game with learning relies on the construction of the following continuum of final-offer games without learning.

Recall that in Section 2 the arbitrator's ideal settlement was denoted by z , and the parties' shared belief about z was denoted by $F(z)$. Suppose instead that for some fixed $k \in (-\infty, \infty)$, the arbitrator's ideal settlement is

$$(27) \quad z_k(s_a) = \frac{(h + h_p)s_a - khm}{(1 - k)h + h_p}.$$

Note that if the ideal settlement is given by (27) then there is nothing the arbitrator needs to learn from the parties: s_a completely determines the arbitrator's ideal settlement, independent of s_p . In this case, the only role for s_p is to help the parties forecast $z_k(s_a)$. Given s_p , the parties' shared belief about $z_k(s_a)$ is Normal with mean $M_k(s_p)$ and precision H_k , where

$$(28) \quad M_k(s_p) = \frac{(1 - k)hm + h_p s_p}{(1 - k)h + h_p} \quad \text{and}$$

$$(29) \quad H_k = \frac{h_a[(1 - k)h + h_p]^2}{(h + h_p)(h + h_p + h_a)}.$$

From (5) and (6), the parties' equilibrium offers in this model of final-offer arbitration without learning are

$$(30) \quad y_{uk}(s_p) = M_k(s_p) + \sqrt{\frac{\pi}{2H_k}} \quad \text{and}$$

$$(31) \quad y_{ek}(s_p) = M_k(s_p) - \sqrt{\frac{\pi}{2H_k}}.$$

These equilibrium offers from the final-offer game without learning will turn out to be identical to the equilibrium offers in this section's model of conventional arbitration with learning!

Imagine that in the model of conventional arbitration with learning, the arbitrator believed that the parties' offer strategies were (30) and (31). The arbitrator could then invert the offer strategies so as to interpret the offers as implicit claims about the parties' information. Given the offer y_u from the union, for instance, the arbitrator could solve (30) for s_p to arrive at the implicit claim that the parties' private information is

$$(32) \quad s_{puk}(y_u) = \sigma_k \left[y_u - \sqrt{\frac{\pi}{2H_k}} \right], \quad \text{where}$$

$$(33) \quad \sigma_k(x) = x + \frac{(1-k)h}{h_p}(x-m)$$

solves $M_k(s_p) = x$ for s_p . Similarly, given the offer y_e from the employer, the arbitrator could solve (31) for s_p to arrive at the implicit claim that the parties' private information is

$$(34) \quad s_{pek}(y_e) = \sigma_k \left[y_e + \sqrt{\frac{\pi}{2H_k}} \right].$$

Note that if the parties' offers are in fact determined by (30) and (31) then these implicit claims will be identical even though the offers are not. If the implicit claims agree, as they will in equilibrium, then it is natural for the arbitrator to believe the common claim. Everything hinges, however, on what the arbitrator believes when the implicit claims disagree.

Suppose that after observing the signal s_a and the offers y_e and y_u the arbitrator arrives at the point estimate $s_p = s_p(y_e, y_u, s_a)$. Then sequential rationality dictates that the arbitrator impose the ideal settlement $M_{pa}[s_p(y_e, y_u, s_a), s_a]$ from (11), or

$$(35) \quad y_a(y_e, y_u, s_a) = \frac{hm + h_p s_p(y_e, y_u, s_a) + h_a s_a}{h + h_p + h_a}.$$

It turns out to be more convenient to express the arbitrator's inference rule not as $s_p(y_e, y_u, s_a)$, but as $M_k(y_e, y_u, s_a)$, where (28) yields

$$(36) \quad M_k(y_e, y_u, s_a) = \frac{(1-k)hm + h_p s_p(y_e, y_u, s_a)}{(1-k)h + h_p}.$$

This last piece of notation makes it possible to state the section's main result in a compact form. (The proof of the Proposition is given in Appendix 2.)

Proposition 3. For each $k \in (-\infty, \infty)$, there exists a separating perfect Bayesian equilibrium of the conventional arbitration game. In each of these equilibria, the arbitrator's ideal (and imposed) settlement is (35). For each fixed k , the parties' offer strategies are (30) and (31), and the arbitrator's inference rule is

$$M_k(y_e, y_u, s_a) = \begin{cases} y_u & \text{if } s_{pek}(y_e) \neq s_{puk}(y_u) \text{ and } z_k(s_a) > \bar{y}, \\ \bar{y} & \text{if } s_{pek}(y_e) = s_{puk}(y_u), \\ y_e & \text{if } s_{pek}(y_e) \neq s_{puk}(y_u) \text{ and } z_k(s_a) \leq \bar{y}. \end{cases}$$

Given k , the arbitrator's ideal settlement can be written

$$y_a = \alpha_k \bar{y}_k + (1-\alpha_k) \frac{khm + h_a s_a}{kh + h_a}, \quad \text{where}$$

$$\bar{y}_k = \frac{(1-k)hm + h_p s_p}{(1-k)h + h_p} \quad \text{and}$$

$$\alpha_k = \frac{(1-k)h + h_p}{h + h_p + h_a}.$$

Several observations about this plethora of equilibria are in order. First, as noted earlier, all of the equilibria described in Proposition 3 are payoff-equivalent: even though changes in k affect both the mean of and the distance between the parties' offers, differences in k are simply differences in the languages used to communicate. Second, two values of k are simple to interpret: $k = 0$ and $k = 1$. When $k = 0$, (28) becomes $M_k(s_p) = M_p(s_p)$, so in equilibrium the arbitrator interprets \bar{y} as the parties' expectation of s given s_p . When $k = 1$, (28) becomes $M_k(s_p) = s_p$, so in equilibrium the arbitrator interprets \bar{y} as s_p itself. And third, in the separating equilibrium with $k = 0$ the parties' offers in conventional arbitration are identical to their offers in the separating equilibrium in final-offer arbitration described in Section 4.

Although the inference rule $M_k(y_c, y_u, s_a)$ stated in the Proposition depends on knife-edge judgments by the arbitrator (i.e., whether $s_{pek}(y_c) = s_{puk}(y_u)$ or not), it has the same spirit as a decision rule that seems reasonable in practice: The arbitrator gives both parties the opportunity to present their cases fully and to challenge assertions in the other's case. Afterwards, when all the cards are supposed to be on the table, the arbitrator asks the parties whether they can now resolve their dispute. If they cannot, then the arbitrator (thinking that reasonable people should by this point agree) conducts an investigation, compares the result to the parties' claims, and throws the book at the party who seems to have made the less reasonable claim. Note that this inference rule has the virtue of supporting a separating equilibrium and so providing the arbitrator with the maximum amount of information in imposing a settlement. Thus, for any k , the decision rule described in Proposition 3 would be an optimal rule if the arbitrator could commit.

6. Relationship of the Learning Models to the Empirical Literature

This section describes the connections between the empirical literature on arbitration and the learning models developed in Sections 4 and 5. As described in the Introduction, much of the empirical work on arbitrator behavior assumes the following model of the arbitrator's ideal settlement:

$$y_a = \alpha \bar{y} + (1-\alpha) y^*(f).$$

As Propositions 2 and 3 emphasize, the arbitrator's equilibrium behavior in this paper takes exactly this form. The paper thus provides a structural interpretation for the empirical literature's reduced-form model: in both conventional and final-offer arbitration, the arbitrator's ideal settlement depends on the parties' offers because the offers convey information, and the importance of the offers (as measured by α) increases as the importance of this information increases relative to the publicly observable information and the arbitrator's independent investigation (i.e., as h_p increases relative to h and h_a).⁴

The paper also fits nicely with two of the stylized facts uncovered by the empirical work. The first stylized fact is that the arbitrator's ideal settlement is identical in the two forms of arbitration; see Ashenfelter-Bloom and Farber-Bazerman. The separating equilibria in the learning models analyzed here are founded on just such an invariant notion of an ideal settlement: under any form of arbitration, if the arbitrator perfectly infers s_p from the parties' offers then the ideal settlement is $M_{pa}(s_p, s_a)$ given in (11).

⁴This structural interpretation is incorrect in the conventional arbitration equilibria in which $\alpha_k > 1$, which occur when $k < -(h_a/h)$, but empirical estimates suggest that $\alpha < 1$.

The second stylized fact is that the parties' offers are more important in the arbitrator's decision-making (i.e., α is higher) when they are closer together; see Farber-Bazerman. The learning models in both Sections 4 and 5 predict that α should be higher when the offers are closer together, but suggest that this need not be a causal relationship. Consider Proposition 2, for instance. Note that α increases with h_p . In equilibrium, the difference between the offers in (25) and (26) is $\sqrt{2\pi/H'}$, where H' is given by (24) and also increases with h_p . Thus, an increase in h_p results in both an increase in α and a decrease in the difference between the offers, thereby producing the observed correlation. The causal story behind the correlation is as follows. The arbitrator puts more weight on the parties' offers when the parties' information is known to be more precise, and hence more valuable in the arbitrator's inference process. At the same time, when the parties' information is more precise, the parties' belief about the arbitrator's private signal, s_a , is more precise, and this causes the parties to submit offers that are closer together.

The model of final-offer arbitration developed in Section 4 does not correspond well to one other stylized fact: Ashenfelter and Bloom find that the union's offer is selected as the settlement more than half the time, whereas in Section 4 the parties' offers are equally likely to be selected. Farber's model of final-offer arbitration without learning is consistent with the Ashenfelter-Bloom finding if the union is more risk-averse than is the employer, and so prefers to submit a conservative offer with a large probability of being chosen rather than an aggressive offer with a small probability. It would be interesting to know whether a learning model with parties of differing risk aversion could integrate the Ashenfelter-Bloom finding on offer-selection frequencies with the role of the parties' offers in the arbitrator's inference process emphasized here.

7. Extensions

Although this paper satisfies three of the conditions that will be required of a complete theory of arbitration, much work remains to be done, both in extending the models analyzed above and in building new models. This section is limited to describing two extensions of the models analyzed above and to providing further discussion of the paper's two most important assumptions.

7a. Comparing Conventional and Final-Offer Arbitration

This paper analyzes conventional and final-offer arbitration within a single economic environment. As explained in the Introduction, this is a prerequisite for a complete theory of arbitration. The obvious next step is a welfare comparison of these and other forms of arbitration. Unfortunately, the model analyzed here is not rich enough to make the comparison interesting. First, because the parties are risk-neutral and have strictly opposed preferences, any preference by one party for one form of arbitration would be exactly reversed for the other party. And second, the arbitrator clearly cannot improve upon the separating equilibrium in conventional arbitration, where the ideal settlement is based on full information and the actual settlement is unconstrained.

A more interesting comparison might arise if the parties were risk-averse. If a separating equilibrium exists in conventional arbitration when the parties are risk-averse then this will again be the arbitrator's preferred form of arbitration. The parties, however, might find that final-offer arbitration offers better insurance than conventional arbitration does. If the parties are allowed to choose the form of arbitration then the arbitrator's preference for conventional arbitration may be overridden.

7b. Modeling Other Forms of Arbitration

Ashenfelter, Dow, and Gallagher (1986) study an interesting variant of final-offer arbitration that is practiced in Iowa. The rules are as follows: first, an arbitrator, acting as a factfinder, suggests a settlement; then the parties either accept the settlement or continue their dispute and submit offers; finally, if the dispute continues then a second arbitrator (drawn from the same pool as the first) chooses one of the parties' offers or the factfinder's proposed settlement.

It is straightforward to extend the model described in Section 4 in order to analyze this version of final-offer arbitration. Suppose the factfinder's investigation yields the signal $s_f = s + \varepsilon_f$, where ε_f is Normal with zero mean and precision h_f and is independent of the other random variables in the model. Then (as in the conventional-arbitration model of Section 5) the factfinder's proposed settlement should be $y_f(s_f) = M_f(s_f)$, given by (9). The rest of the model is as before, except for two changes. First, both parties and the second arbitrator update their beliefs after observing the factfinder's proposal. And second, the second arbitrator now has a third choice in the feasible set of imposed settlements.

Ashenfelter, Dow, and Gallagher argue that the "extent to which the second arbitrator does not concur in the factfinder's proposal is an obvious measure of the degree of arbitral uncertainty that exists in the system" (p. 2). While this interpretation is certainly plausible, the disagreement between arbitrators also may reflect the second arbitrator's learning from the parties' offers.

7c. Commitment by the Arbitrator

One of the two important assumptions in this paper is that the arbitrator is unable to commit ex-ante to a decision rule, even though this might induce a socially desirable

outcome of the negotiation and arbitration process. This assumption seems to represent observed practice accurately: arbitrators do not commit to decision rules before the parties make their offers. On the other hand, the assumption ignores the important possibility that an arbitrator might develop a reputation in the course of repeated play. The arbitrator-selection mechanisms described and analyzed by Bloom and Cavanagh (1986a,b) will play an important role in the eventual analysis of such dynamic games. Because the complete dynamic analysis is likely to be difficult, however, it would be worthwhile to analyze the models that complement this paper---static models in which the arbitrator is allowed to commit to a decision rule *ex-ante*. As noted in Section 5, each of the continuum of decision rules described there is an optimal rule if the arbitrator can commit in conventional arbitration. The nature of the analogous optimal rule under commitment in final-offer arbitration is an open question.

7d. Arbitration Between Parties with Symmetric Information

The second important assumption in this paper is that the parties have symmetric information about the state of the employment relationship. It seems likely, instead, that each of the parties will have private information about some aspects of the employment relationship. Such an information asymmetry between the parties may influence both the means and the substance of the parties' equilibrium communication with the arbitrator. Perhaps more importantly, it prompts the question of how the parties came to be in arbitration.

The theory of bargaining under symmetric information, due to Nash (1950) and Rubinstein (1982), suggests that negotiations should not fail when information is symmetric, especially when the arbitration game that follows failed negotiations imposes

costs (such as the arbitrator's fee) on the parties. This does not imply, however, that the prospect of arbitration has no effect. Crawford (1982a), for instance, analyzes the way arbitration changes the threat point in a Nash bargaining game: arbitration has an effect even though the bargainers never fail to settle. Crawford's analysis involves models of conventional and final-offer arbitration without learning. The analogous analysis using this paper's learning models seems likely to imply that the prospect of arbitration has a smaller effect on negotiated settlements when learning is more important, because learning links the arbitrated settlement to the state and so prevents a disadvantaged bargainer from escaping an unattractive negotiated settlement through arbitration by an uninformed arbitrator.

Thus, the prospect of arbitration between parties with symmetric information can have an effect even if the parties never enter arbitration. Furthermore, in richer models than those developed by Nash and Rubinstein, disputes can occur between parties with symmetric information. Crawford (1982b), for instance, formalizes Schelling's (1956) view that bargaining is a struggle between the bargainers to commit themselves to favorable positions. Attempts at commitment are uncertain to succeed, but are irreversible if they do. If only one party's attempt succeeds then that party wins a favorable negotiated settlement, but if the parties are both committed to incompatible positions then negotiations fail and arbitration may occur. A model based on these ideas can produce disputes between parties with symmetric information, but more work needs to be done in modeling the process of commitment.

One intriguing approach to modeling the process of commitment emphasizes the idea that the parties are groups rather than individuals; see Walton and McKersie (1965). A union leader, for example, can try to incite and yet control the rank and file, but internal politics and group dynamics may make the rank and file's reactions beyond the union

leader's control. In this case, the union leader may become committed in the sense envisioned by Schelling and Crawford. Alternatively, when the parties are groups a dispute may occur because of asymmetric information between them about features of their relationship that are of no concern to the arbitrator, including perhaps the career concerns of individual negotiators. In this case, with regard to the learning in the arbitrator's inference problem the parties are appropriately viewed as having symmetric information, even though asymmetric information on another dimension caused the dispute. These possibilities suggest that it may be a mistake to take too literally models that approximate a firm and a union as single-agent bargainers.

APPENDIX 1

This appendix establishes the global optimality of the offers that satisfy the first-order conditions (5) and (6), (16) and (17), (25) and (26), and (30) and (31). An explicit proof is given for the union's offer specified in (5). All the other cases are analogous.

Let the union's expected payoff given in (1) be denoted

$$(A1) \quad U(y_u | y_e) = y_e F(\bar{y}) + y_u [1 - F(\bar{y})],$$

where $\bar{y} = (y_u + y_e)/2$ is the average of the offers and F is a Normal distribution with mean M and precision H . In this notation, the first-order condition (3) is simply $U'(y_u | y_e) = 0$, where the prime denotes a derivative with respect to y_u , and

$$(A2) \quad U'(y_u | y_e) = [1 - F(\bar{y})] - \frac{1}{2}(y_u - y_e)f(\bar{y}).$$

The goal of this appendix is to prove that $y_u = M + \sqrt{\pi/2H}$ achieves the global maximum of $U(y_u | y_e^*)$ where $y_e^* = M - \sqrt{\pi/2H}$ from (6).

It suffices to prove that U is quasi-concave in y_u when $y_e = y_e^*$. This is done in two parts. The first establishes that U is concave for $y_u \in [y_e^*, M + D]$, where $D > 0$. And the second establishes that $U' < 0$ for $y_u \in (M + d, \infty)$, where $0 < d < D$. (Recall that union offers below the employer's offer are sub-optimal for the union, and so need not be considered here.)

The concavity of U on $[y_e^*, M + D]$ follows from evaluating U'' on this interval:

$$(A3) \quad U''(y_u | y_e^*) = -f(\bar{y}^*) - \frac{1}{4}(y_u - y_e^*)f'(\bar{y}^*),$$

where $\bar{y}^* = (y_u + y_e^*) / 2$. For the Normal distribution with mean M and precision H , $-f'(\bar{y}^*) = H(\bar{y}^* - M)f(\bar{y}^*)$, so

$$(A4) \quad U''(y_u | y_e^*) = f(\bar{y}^*) \left\{ -1 + \frac{H}{4}(y_u - y_e^*)(\bar{y}^* - M) \right\},$$

which is negative for $(y_u - M)^2 < \frac{1}{H} \left(8 + \frac{\pi}{2} \right) = D^2$.

The negative slope of U on $(M + d, \infty)$ follows from a simple inequality:

$$(A5) \quad H(\bar{y}^* - M) [1 - F(\bar{y}^*)] < f(\bar{y}^*).$$

This inequality follows from integrating $(x - \bar{y}^*) f(x)$ over (\bar{y}^*, ∞) when x is Normal with mean M and precision H . Given (A2) and (A5), it suffices to show that

$$(A6) \quad (y_u - y_e^*)(\bar{y}^* - M) > \frac{2}{H},$$

which holds for $(y_u - M)^2 > \frac{1}{H} \left(4 + \frac{\pi}{2} \right) = d^2$.

APPENDIX 2

This appendix proves Proposition 3. The proof that the arbitrator's strategy is a sequentially rational best response to the parties' strategies is simple. Using (33), the inference rule $M_k(y_e, y_u, s_a)$ given in the Proposition can be expressed as the inference rule $s_p(y_e, y_u, s_a)$ used in (35). This inference rule is clearly correct in equilibrium and is well defined off the equilibrium path; given this belief, (35) guarantees sequential rationality.

The proof that the parties' strategies are best responses to each other and to the arbitrator's strategy is slightly more involved. The central point is that out of equilibrium in the conventional-arbitration game, it is as though the arbitrator makes the parties play a final-offer game that has a unique Nash equilibrium that yields both parties the expected payoff $M_k(s_p)$, which equals their equilibrium expected payoff in conventional arbitration. Consider the payoff to, say, the employer from submitting an offer different from the equilibrium offer given by (31). (A deviation by the union is analyzed symmetrically.) By hypothesis, the union's offer will be given by (30), so any deviation by the employer will cause $s_{pek}(y_e) \neq s_{puk}(y_u)$. This in turn causes the arbitrator to use one of the two out-of-equilibrium parts of the inference rule $M_k(y_e, y_u, s_a)$ specified in the Proposition. Therefore, the problem of finding the optimal misleading claim, s_{pek} , for the employer to submit in the conventional-arbitration game is equivalent to the problem of finding the optimal offer, y_e , for the employer to submit in the final-offer game analyzed in the text. But the solution to this second problem is given by (31) and yields a payoff identical to the employer's equilibrium payoff in the conventional-arbitration game. Thus, no deviation is profitable.

REFERENCES

- Ashenfelter, O. and D. Bloom, "Models of Arbitrator Behavior: Theory and Evidence," American Economic Review, 74 (1984), 111-24.
- _____, J. Dow, and D. Gallagher, "Arbitrator and Negotiation Behavior under an Appellate System," Princeton mimeo, August 1986.
- Bloom, D., "Empirical Models of Arbitrator Behavior Under Conventional Arbitration," Review of Economics and Statistics, 68 (1986), 578-85.
- _____ and C. Cavanagh, "An Analysis of the Selection of Arbitrators," American Economic Review, 76 (1986a), 408-22.
- _____ and _____, "An Analysis of Alternative Mechanisms for Selecting Arbitrators," Harvard Institute of Economic Research D.P. # 1224, April 1986b.
- Crawford, V., "Compulsory Arbitration, Arbitral Risk, and Negotiated Settlements: A Case Study in Bargaining Under Imperfect Information," Review of Economic Studies, 49 (1982a), 69-82.
- _____, "A Theory of Disagreement in Bargaining," Econometrica, 50 (1982b), 607-637.
- _____ and J. Sobel, "Strategic Information Transmission," Econometrica, 50 (1982), 1431-1451.
- DeGroot, M., Optimal Statistical Decisions, New York: McGraw-Hill, 1970.
- Farber, H., "An Analysis of Final-Offer Arbitration," Journal of Conflict Resolution, 35 (1980), 683-705.

- _____ and M. Bazerman, "The General Basis of Arbitrator Behavior: An Empirical Analysis of Conventional and Final-Offer Arbitration," Econometrica, 54 (1986), 819-44.
- Grossman, S. and O. Hart, "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," Journal of Political Economy, 94 (1986), 691-719.
- Kalai, E. and R. Rosenthal, "Arbitration of Two-Party Disputes Under Ignorance," International Journal of Game Theory, 7 (1979), 65-72.
- Nash, J., "The Bargaining Problem," Econometrica, 18 (1950), 155-162.
- Rubinstein, A., "Perfect Equilibrium in a Bargaining Model," Econometrica, 50 (1982), 97-109.
- Schelling, T., "An Essay on Bargaining," American Economic Review, 46 (1956), 557-73.
- Stevens, C., "Is Compulsory Arbitration Compatible with Bargaining?," Industrial Relations, 5 (1966), 38-52.
- Walton, R. and R. McKersie, A Behavioral Theory of Labor Negotiations. New York: McGraw-Hill Co., 1965.
- Williamson, O., Markets and Hierarchies. New York: The Free Press, 1975.