

Career Concerns of Bargainers

John Fingleton
Irish Competition Authority

Michael Raith
University of Rochester and CEPR

This article studies strategic bargaining in which a seller and a buyer are each represented by an agent. Potential agents differ in their ability to obtain information about the other party's reservation price; neither principal knows the other's reservation price or her agent's type. Agents are motivated by career concerns; they want to be perceived as skilled bargainers by their principals. In equilibrium, skilled agents use their private information optimally, while unskilled agents randomize between aggressive and soft price bids, attempting to imitate skilled types. We compare "open-door" bargaining, in which principals can observe the entire bargaining game as well as its outcome, with "closed-door" bargaining, in which they observe only the outcome. We show that agents unambiguously bargain more aggressively with open doors than behind closed doors, which leads to a less efficient bargaining outcome. Their principals may therefore prefer to let their agents bargain behind closed doors.

1. Introduction

In most bargaining situations, negotiators bargain not for themselves, but on behalf of others. Managers negotiate on behalf of shareholders, union leaders on behalf of workers, lawyers on behalf of clients, politicians on behalf of voters, diplomats on behalf of countries (see Mnookin and Cohen, 1999). While bargaining theory has helped us understand many patterns of bargaining between players negotiating in their own interest, little is known about how the agency relationships between negotiators and their principals impact on bargaining behavior and outcomes.

We benefited from many conversations with Canice Prendergast, and also thank two anonymous referees, Judy Chevalier, Wouter Dessein, Clemens Esser, Guido Friebel, Luis Garicano, Paul Heidhues, Randy Kroszner, Allison Raith, Matthias Raith, Philip Reny, Lars Stole, and seminar participants at the University of Chicago Graduate School of Business, London School of Economics, University of Michigan, Stockholm School of Economics, and at several conferences for helpful comments and suggestions. Any errors are our own. Thanks also to the Graduate School of Business at the University of Chicago for hosting John Fingleton as a visiting scholar and the European Centre for Advanced Research in Economics (ECARE) in Brussels where this research was initiated.

The Journal of Law, Economics, & Organization, Vol. 21, No. 1,
doi:10.1093/jleo/ewi008

© The Author 2005. Published by Oxford University Press. All rights reserved. For Permissions, please email: journals.permissions@oupjournals.org

In this article we address a question of particular practical relevance, namely to what extent principals ought to monitor agents who bargain on their behalf. Standard principal-agent theory suggests that monitoring mitigates agency problems and hence is unambiguously beneficial, implying that monitoring should be limited only by its direct costs.

Yet among practitioners, there is a strong belief that monitoring often has an adverse effect on the behavior of agents. For example, Fisher and Ury (1981:36) recommend: "To reduce the dominating and distracting effect that the press, home audiences, and third parties may have, it is useful to establish private and confidential means of communicating with the other side" (see also King and Zeckhauser, 1999; Kurtzberg et al., 1999). Similarly, in 1995 the Court of First Instance of the European Communities (1995) rejected demands for the disclosure of minutes of meetings of the European Council, arguing that "The process of compromise and negotiation . . . would be jeopardized if delegations were constantly mindful of the fact that the positions they were taking . . . could at any time be made public." More often than not, negotiations in business, legal disputes, politics, and diplomacy appear to be surrounded by secrecy.

This article offers an explanation for why monitoring bargaining agents can be inefficient. We study a model in which monitoring an agent's actions (which we refer to as "open-door bargaining") induces the agent to bargain more aggressively, that is, to engage in "posturing."¹ If this occurs on both sides of the bargaining table, delay or disagreement may follow. Principals may therefore prefer to let agents negotiate "behind closed doors," that is, to learn the outcome of bargaining but not the agents' actions that led to it. Although in general terms, this idea is familiar and can be traced back at least to Schelling (1960), our article is the first to give it a rigorous theoretical foundation. We identify the conditions that lead to posturing and (in the Conclusion) discuss a number of practical implications for the principals' choice between open- and closed-door bargaining. We provide sufficient conditions for when delegating bargaining to an agent is optimal; however, our focus is on the consequences of monitoring, taking delegation as given.

A central assumption in our theory is that agents are motivated by career concerns rather than explicit incentives. Agents often have little direct stake in the outcome of bargaining, and offering them explicit pay for performance is impractical in complex bargaining situations with multiple dimensions. Instead, agents seek to be perceived as good bargainers in order to be reappointed by their principals.²

1. This effect has been confirmed in a number of experiments and field studies; cf. Rubin and Brown (1975:43–54) and Pruitt (1981:41–45).

2. Conversely, when incentive contracts are impractical, the best principals can do is to hire agents who have performed well before, as Francis Bacon recommended in his essay *Of Negotiating*: "Use such persons . . . as have benee luckie, and prevailed before in things wherein you have employed them; for that breeds confidence, and they will strive to maintaine their prescription" (Bacon, 1862 [1625]:196).

We study bargaining between two parties (i.e., principals) that are incompletely informed about each other's preferences. We assume that some agents are more skilled than others at evaluating or obtaining information about the other party. On average, they negotiate better deals for their principals, and are in this sense better bargainers. Principals draw inferences about their agents from the observed bargaining process. Agents, in turn, choose their strategies to promote their reputation as skilled bargainers.

In the simplest setup of our model (see Section 2), a seller bargains through an agent with a buyer. The agent and the buyer play a one-shot bargaining game in which one side is randomly chosen to make a price offer and the other side accepts or rejects. The seller and buyer do not know each other's reservation price, and with positive probability, trade is not efficient. The agent may be skilled, in which case he receives an informative signal about the buyer's reservation price, or unskilled, in which case his signal is uninformative. The agent knows whether he is skilled, but the seller does not know her agent's type. We compare two scenarios: *open-door bargaining*, where the seller observes the entire bargaining game, and *closed-door bargaining*, where she observes only the final outcome, that is, trade at some price or disagreement.

An agent's *reputation* is the posterior probability, as evaluated by the seller, that the agent is skilled. We assume that the agent cares only about his reputation. Price bids are cheap talk; prices and outcomes affect the agent's utility only insofar as they affect his reputation. Thus, in a technical sense, our article is similar to recent "expert" models in which agents attempt to signal that they have high-quality information about an unknown state of the world (cf. our discussion of the literature below).

The nature of equilibrium can best be understood by focusing on the seller's side, taking the buyer's strategy as given. An agent of a seller with a low reservation price must, when making an offer, decide whether to bid a high or a low price. A low price is always accepted by the buyer, but results in a low payoff for the seller. A high price is accepted if the buyer's reservation price is high, but rejected if it is low, in which case bargaining ends with disagreement, the worst outcome. The agent thus faces a choice between the safe option of demanding a low price and the risky option of demanding a high price.

With both open and closed doors, there exists a unique perfect Bayesian equilibrium in which the seller's expected payoff is increasing in the agent's skill (cf. Section 3). A seller with a low reservation price believes that her agent is more likely to be skilled if he sells at a high price, less likely if he sells at a low price, and least likely if no trade occurs. Given the seller's beliefs, a skilled agent makes efficient use of his information; that is, he demands a high price if he believes that the buyer's reservation price is high, and a low price otherwise. An unskilled agent, in contrast, randomizes between a low and a high price, and thus mimics the behavior of a skilled agent to some extent. As a result, a skilled agent is more likely to sell at a high price and to avoid disagreement, which confirms the principal's beliefs.

An important feature of the equilibrium is that with open or closed doors, an unskilled agent demands a high price (i.e., chooses the risky strategy) less often

than does a skilled agent. To see why, suppose that an unskilled agent demanded a high price more often. Then the expected reputation from demanding a low price, which is the same for both skilled and unskilled agents, would exceed the prior probability that an agent is skilled. That, however, would mean that both skilled and unskilled agents could secure an expected reputation better than the prior, which is not possible.

Our main result is that with open-door bargaining, an unskilled agent bargains more aggressively, that is, bids a high price more often, than with closed doors (Section 4). Since for any given strategy of an unskilled agent, open-door bargaining is more informative for the principal than closed-door bargaining because she can observe which side makes an offer, it is optimal for an unskilled agent to mimic a skilled agent to a greater extent than with close doors. Since in equilibrium an unskilled agent demands a high price less often than does a skilled agent, mimicking a skilled agent to a greater extent means demanding a high price more often than with closed doors. The open-door equilibrium is less efficient because disagreement occurs more often.³

This result does not depend on the particular structure of the bargaining game; it is a general feature of a model in which agents must choose between a safe and a risky strategy. Two assumptions are important, however. The first is that the agent knows his level of skill. Without this knowledge, the agent's strategy with open doors tends to be the same as with closed doors. The second assumption is that the agent maximizes his reputation with the seller, who knows her own reservation price. If, instead, the agent maximizes his reputation on an external market that does not know either side's reservation price, there is no systematic difference in behavior between open and closed doors.

We have so far taken the buyer's bargaining behavior as given. When the buyer bargains for herself, then in our one-shot model, her optimal strategy does not depend on the agent's strategy, and hence does not depend on whether bargaining is with open or closed doors.⁴ If the buyer is also represented by an agent with reputation concerns (Section 5.2), open-door bargaining is unambiguously less efficient than closed-door bargaining because unskilled agents on both sides will bargain more aggressively. The seller and buyer may still prefer open-door bargaining if the gain in information about their agents is more valuable than the short-term efficiency loss. However, such a gain may not exist if the distortion caused by more aggressive bargaining is large.

One may be tempted to think that open-door bargaining emerges as the outcome of a Prisoners'-dilemma situation in which each side unilaterally has an incentive to monitor the agents. That is not necessarily the case (Section 5.3): unilateral monitoring is unlikely to be profitable because the own agent's more

3. Since our bargaining model is static, and since trade is always efficient when it occurs, disagreement when trade would be efficient is the only source of inefficiency. As we argue in Section 4, however, the intuition for our main result is likely to be quite general. In a dynamic bargaining situation, for instance, an inefficient bargaining outcome may also occur because of delay.

4. If bargaining proceeds in multiple rounds, one might expect the buyer to respond to a more aggressive strategy employed by the seller's agent by making concessions more easily, see Section 5.1. Open-door bargaining might then be a means for the seller to gain an advantage over the buyer.

aggressive bargaining behavior may prompt the other agent to be more aggressive too. That is, an agent with reputation concerns reacts differently from a bargainer who maximizes the payoff from trade.

Three articles related to ours are Perry and Samuelson (1994), Cai (2000), and Prat (2003). In Perry and Samuelson (1994), heterogeneous constituents bargain through an agent with another party. With open-door but not closed-door bargaining, the constituents can terminate bargaining after initial offers have been made and rejected. This may occur if current constituents fear that due to changes in the population, future constituencies might approve of an agreement that makes the surviving current constituents worse off. The threat of termination induces the other party to make a larger concession in the first round, implying that open-door bargaining is advantageous for the constituents.

In Cai (2000), too, constituents contract with an agent to bargain with another party. Potential agents differ in their private costs of bargaining; the constituents face uncertainty about both their agent's bargaining costs and the bargaining surplus at stake. In equilibrium, low-cost agents signal their type by excessively delaying negotiations, which is inefficient. Cai does not compare different monitoring structures, but instead focuses on the loss of efficiency resulting from delegation. More importantly, Cai rules out that a principal can compensate her agent for his costs of bargaining. In our model, in contrast, agents differ in their ability to ascertain the other party's reservation price, and it is impossible for principals to pay unskilled agents to be skilled.

Finally, our article is in a technical sense related to recent "expert" models in which agents maximize their reputation for having precise information about an unknown state of the world (cf. Scharfstein and Stein, 1990; Prendergast and Stole, 1996; Ottaviani and Sørensen, 2000; and Prat, 2003). Among these, only Prat is concerned with a principal's optimal choice of information structure. Prat shows that if an agent does not know his type, and if the signal he receives is sufficiently highly correlated with his type, the principal may be better off if she can observe only the consequence of the agent's action, but not the action itself. If the agent does know his type, the resulting equilibrium is similar to the one obtained here, but in contrast to our main result, there is no systematic difference between the cases of observable versus unobservable actions.

The reason for this difference is that Prat's model and most other expert models are symmetric; states, actions, and outcomes are not ranked in any way. Here, in contrast, the agent's choice between a risky and a safe action leads to an equilibrium in which some outcomes are "good" for the agent's reputation and others "bad," irrespective of the true state of the world. This is an altogether very different game compared with one in which the agent's objective is simply to correctly guess the state of the world.

This article is the first to provide a rigorous foundation for predictions that are both intuitive and in agreement with practitioners' views on the drawbacks of open-door bargaining. In particular, both moral hazard models and Holmström's

(1999) career-concerns model fail to explain how inefficient posturing can arise. The problem is that in bargaining with incomplete information, there is no monotonic relationship between an agent's effort and the resulting price bids, or between price bids and the resulting payoff or reputation. For example, monitoring might induce an agent to work harder, for example, to learn the other side's reservation price. Better information, however, would not systematically lead to more aggressive bargaining; on the contrary, bargaining would become more efficient.

A different argument, namely that monitoring allows a principal to *commit* to a bargaining strategy through a contract with her agent, runs into two problems. First, without a genuine agency problem, commitments through contracts are not credible if they can be undone by secret renegotiation (cf. Katz, 1991). The second problem is the above-mentioned nonmonotonicity between prices and payoffs when information is incomplete: even if a principal could credibly commit herself, she would avoid adopting a too aggressive bargaining position because of the risk of rejection. In fact, the results of Myerson and Satterthwaite (1983) and Cramton (1985) suggest that if *both* sides can commit to prices rather than make gradual concessions over time, the resulting bargaining outcome would be more efficient, not less (cf. also Crawford, 1982). Inefficient posturing, as observed in practice, therefore cannot easily be explained by the use of commitment tactics.

2. Model

A seller S employs an agent A_s to bargain with a buyer B over an indivisible good (in Section 5, we consider the case where B herself is represented by an agent). The seller's reservation price is either \underline{s} or \bar{s} , and the buyer's reservation price is either \bar{b} or \underline{b} , where $\underline{s} < \underline{b} < \bar{s} < \bar{b}$ (see Figure 1). We call a seller with valuation \underline{s} and a buyer with valuation \bar{b} "weak," and the other types "tough." Both S and B are weak with probability p and tough with probability $1 - p$; the reservation prices of S and B are independent random variables. Since $\underline{b} < \bar{s}$, trade is inefficient with positive probability. These assumptions are similar to those of Chatterjee and Samuelson (1987).

Bargaining is conducted through a one-shot game in which, with equal probability, either A_s or B is chosen to make an offer. Two different prices can be submitted, either high (H) or low (L), where $\underline{s} < L < \underline{b} < \bar{s} \leq H < \bar{b}$ (see Figure 1). The other player can accept the offer, in which case trade takes place at that price, or reject it, in which case no trade takes place and both S and B earn zero payoff. The agent may not submit a bid that can lead to a loss for S , but is otherwise not restricted in the choice of his bid. Thus if S is weak, A_s can offer or accept L or H , whereas if S is tough, he can offer or accept only H .⁵ We will say that A_s plays tough if he bids H , and weak if he bids L .

5. In Section 3 we derive a sufficient condition for when it is indeed optimal for a weak seller to give her agent discretion instead of requiring him always to bid one particular price.




Figure 1. Valuations and prices

To focus on the principal-agent relationship between S and A_s , we take, for now, the buyer's bargaining behavior as given and assume that a weak B bids L (i.e., makes a tough bid) with probability a_B (a tough B always bids L). In Section 5 we endogenize the buyer's bargaining behavior.

The seller hires A_s from a pool of candidates who are "good" with probability θ and "bad" with probability $1 - \theta$; S cannot observe A_s 's type. A good agent receives a signal about B 's reservation price that is correct with probability $q \in (1/2, 1)$; a bad agent does not receive any information⁶. We assume that A_s knows his type, but cannot credibly communicate it to S . The motivation for this assumption is that while any agent can form an opinion about the other bargaining party's preferences, agents typically also have private information about how well-founded their opinion is.

The seller can choose between two information structures: with *open-door bargaining*, S observes the entire bargaining game. That is, she observes whether A_s or B makes an offer, as well as the offer itself and the other player's response. With *closed-door bargaining*, S observes only the final outcome, that is, the transaction price or disagreement.

Two objectives determine S 's choice between these information structures. One is to maximize her expected payoff from bargaining. Another is to obtain information about her agent's ability in order to compensate him accordingly in future (unmodeled) bargaining situations. We call the probability with which S believes that A_s is good the agent's "reputation." The seller combines the agent's initial reputation θ with information gleaned from observed bargaining behavior to form a posterior reputation $\hat{\theta}$.

Following Holmström (1999) and others, (cf. Scharfstein and Stein, 1990; Ottaviani and Sørensen, 2000; and Prat, 2003), we assume that the agent's sole objective is to maximize his reputation $\hat{\theta}$. In particular, we assume that he cannot be compensated conditional on the outcome of bargaining, nor does he have any intrinsic interest in the outcome. This assumption seems quite realistic: with the exception of litigation lawyers and real estate or athlete's agents, negotiators are typically rewarded for good performance by reappointment, not by explicit pay for performance, because real-world bargaining usually covers multiple issues that often defy measurement.

The timing of the game is as follows:

1. S chooses an information structure: open or closed door.
2. S hires A_s . The agent knows his own type but S does not.
3. The valuations of S and B are realized.

6. These assumptions are very similar to those of Scharfstein and Stein (1990).

4. A_s learns S 's valuation. A good agent also receives a signal about B 's valuation.
5. Bargaining: with probability $1/2$, either A_s or B is chosen to offer a price (L or H). The other player accepts or rejects.

The bargaining model used here is (for reasons of tractability) very simple and specific, but captures the two features of real-world bargaining that are central to the question studied in this article. First, an agent must choose whether to adopt an aggressive or an appeasing bargaining strategy without knowing the other side's reservation price. Appeasing is a safe strategy, but it does not lead to a very favorable agreement. Being aggressive is a risky strategy that may lead to a high payoff, or it may lead to delay or disagreement. Second, closed-door bargaining is strictly less informative for a principal than open-door bargaining. It would be easy to construct alternative models in which the open- and closed-door equilibria are identical. However, since the purpose of our analysis is to understand why in real-world bargaining agents appear to behave differently with open and closed doors, we chose to study a model in which a meaningful difference between the two scenarios exists.

The preceding argument also relates to our restriction of the bargainers' actions to two prices. One might object, for example, that if prices could be freely chosen, then the seller's agent would restrict his price offers to \bar{b} and \underline{b} instead of H and L . Likewise, the buyer would restrict her prices to \bar{s} and \underline{s} . With these choices, the difference between open- and closed-door bargaining would disappear in our model because, even with closed doors, the seller could then infer whether her agent or the buyer made an offer. This alternative setup of the model, however, would fail to capture the fact that in real-world bargaining, principals have strictly less information about the bargainers' actions when bargaining is conducted behind closed doors than with open doors. Thus in a model in which the the players can have only two types, a meaningful difference between our two scenarios can exist only if the feasible prices are restricted as well.⁷

3. Equilibrium

Our model has features of both a cheap-talk game and a signaling game: first, bids are costless and affect A_s 's payoff only indirectly through his principal's inferences. The game therefore has many equilibria; in particular, any pooling strategy is an equilibrium strategy. Second, a good and a bad agent both want to be perceived as good bargainers, that is, they have the same preferences.

7. On the other hand, if the players' types are continuous, then a meaningful difference between open and closed doors exists even with a continuous-type space. For example, for uniformly distributed types and a perfectly informed good agent, it is possible to compute the open-door equilibrium. In that equilibrium, a good agent bids the buyer's reservation price, while a bad agent draws a price from a continuous distribution. All qualitative features of this equilibrium are the same as in our two-type model (see, e.g., footnote 9). Unfortunately, however, it is not possible to compute the closed-door equilibrium, which is why we chose the simpler setup presented here.

Separation in equilibrium is possible because the agents differ in the quality of their information about B . The career-concerns setting suggests a very natural equilibrium selection criterion: S will be willing to pay A_s according to his perceived ability only if the expected payoff from hiring a good agent exceeds the expected payoff from hiring a bad one.⁸ We show that there exists a unique equilibrium with this property, and for the remainder of the article we confine our attention to this equilibrium.

In the following, we restrict our attention to the case of a weak S . If S is tough, A_s can bid or accept only a high price, which means that S does not learn anything about her agent. Thus the possibility that S may be tough will be important later when we determine B 's bargaining strategy, but is irrelevant for S or A_s .⁹ Moreover, when studying A_s 's behavior, we can restrict attention to the case when he makes an offer, since, as argued above, he will always accept any offer made by B .¹⁰

This restriction of relevant cases implies that under each information structure, there are only three events that lead to nontrivial inferences. With closed doors, the observable events are trade at a high price, trade at a low price, and no trade, which we denote by H , L , and N . With open doors, the relevant events are those where A_s makes an offer, which we can again denote by H , L , and N . Thus, in both cases, the outcome of bargaining can be denoted by an event $e \in \{H, L, N\}$.¹¹

Technically there are three types of agents: a good agent with a signal that B is tough ($k = t$), a good agent with a signal that B is weak ($k = w$), and a bad agent ($k = b$). Allowing for mixed strategies, let a_s^k be the probability that a type k agent plays tough. The agent's strategy is then described by the triple $\mathbf{a}_s = (a_s^t, a_s^w, a_s^b)$, and let $\mathbf{a} = (\mathbf{a}_s, a_B)$.¹²

Let $\theta(e|\mathbf{a})$ be S 's posterior probability that A_s is good, which depends on the observed event e , the agent's strategy \mathbf{a}_s , and B 's bargaining behavior a_B . If A_s bids L , then B always accepts, and the resulting reputation for A_s is $\hat{\theta}(L|\mathbf{a})$. From the perspective of a type k agent, the expected payoff from bidding H depends on the subjective probability that B is weak, which we denote by p^k for $k = t, w, b$. For a good agent who receives a signal that B is weak, the conditional probability that B is indeed weak is $p^w = pq/[pq + (1 - p)(1 - q)]$,

8. This is the approach used in all related papers; see Scharfstein and Stein, 1990; Ottaviani and Sørensen, 2000; and Prat, 2003.

9. Similarly, it can be shown that if the distribution of reservation prices were continuous, then the higher S 's reservation price, the less prices A_s would have to choose from, and hence the less informative his actions and the resulting outcome would be.

10. In Fingleton and Raith (2000), we also study a two-stage version of the game described here. In that case, second-round offers are always accepted if compatible with the principal's reservation price, but first-round offers may be rejected. Our results are the same as for the one-stage game.

11. While we use H and L both for price bids and observable events, there should not be any reason for confusion. Similarly, while the events H and L are strictly speaking different for open and closed doors (e.g., the event H with closed-door bargaining is the union of the event H under open-door bargaining and the event where B offers H and A_s accepts), it will always be clear which event is meant.

12. In what follows, a subscript b stands for "buyer" and a superscript b stands for a "bad" agent.

where the second term in the denominator is the probability that B is tough and A_s receives a wrong signal. Similarly, for a good agent who receives a signal that B is tough, the conditional probability that B is weak is $p^t = p(1 - q)/[p(1 - q) + (1 - p)q]$. For a bad agent, we simply have $p^b = p$. The expected payoff from bidding H is then

$$p^k \hat{\theta}(H|\mathbf{a}) + (1 - p^k) \hat{\theta}(N|\mathbf{a}).$$

For a given a_B , a *perfect Bayesian equilibrium* consists of a strategy \mathbf{a}_s and beliefs $\hat{\theta}(e|\mathbf{a})$ such that (1) the action a_s^k is optimal for a type k agent, given S 's beliefs, that is

$$a_s^k = \arg \max_{a \in [0,1]} a[p^k \hat{\theta}(H|\mathbf{a}_s^k(\mathbf{a})) + (1 - p^k) \hat{\theta}(N|\mathbf{a}_s^k(\mathbf{a}))] + (1 - a) \hat{\theta}(L|\mathbf{a}_s^k(\mathbf{a})) \text{ for } k = t, w, b, \tag{1}$$

where $\mathbf{a}_s^k(\mathbf{a})$ denotes the strategy combination \mathbf{a} , with a_s^k replaced by a ; and (2) the seller's beliefs $\hat{\theta}(e|\mathbf{a})$ satisfy Bayes' rule for any event e that occurs with positive probability, that is

$$\hat{\theta}(e) = \frac{\Pr(e|\text{good})\theta}{\Pr(e|\text{good})\theta + \Pr(e|\text{bad})(1 - \theta)}, \tag{2}$$

where $\Pr(e|\cdot)$ is the probability that a good or bad agent reaches event e , evaluated from S 's point of view.

We then have

Proposition 1. With both open- and closed-door bargaining, there exists a unique perfect Bayesian equilibrium satisfying the requirement that S 's expected payoff is higher if A_s is good than if he is bad. In this equilibrium, (i) a good agent who believes that B is weak plays tough ($a_s^w = 1$), (ii) a good agent who believes that B is tough plays weak ($a_s^t = 0$), and (iii) a bad agent plays tough with probability a_s^b , where $0 \leq a_s^b \leq pq + (1 - p)(1 - q)$. The corresponding posterior beliefs are given by

$$\hat{\theta}(H) = \frac{[i(1 - a_B) + a_s^b]\theta}{i(1 - a_B) + [qa_s^w + (1 - q)a_s^t]\theta + (1 - \theta)a_s^b}, \tag{3}$$

$$\hat{\theta}(L) = \frac{[(1 - p + pa_B)i + 1 - (1 - p - q + 2pq)a_s^w - (p + q - 2pq)a_s^t]\theta}{(1 - p + pa_B)i + [1 - (1 - p - q + 2pq)a_s^w - (p + q - 2pq)a_s^t]\theta + (1 - a_s^b)(1 - \theta)}, \tag{4}$$

$$\hat{\theta}(N) = \frac{[qa_s^t + (1 - q)a_s^w]\theta}{[qa_s^t + (1 - q)a_s^w]\theta + a_s^b(1 - \theta)}, \tag{5}$$

where $i = 0$ for open-door and $i = 1$ for closed-door bargaining. In both scenarios, the beliefs satisfy $\hat{\theta}(H) > \theta > \hat{\theta}(L) > \hat{\theta}(N)$.

All proofs are in the appendix. The intuition is the same for both open- and closed-door bargaining. Suppose $\hat{\theta}(e)$ is increasing in the price at which A_s trades, that is, $\hat{\theta}(H) > \hat{\theta}(L) > \hat{\theta}(N)$. Then, since $p^w > p^b > p^t$, the linearity of Equation (1) in p^k implies that at most one type of player uses a mixed strategy. Moreover, a good agent never mixes: if, for instance, a good agent who believed B is weak mixed between L and H , then both a bad agent and a good agent who believed B is tough would play L . But then S would infer that A_s is good whenever H or N occur, in which case A_s would strictly prefer to play H . A similar argument applies for a type t agent. It follows that in equilibrium, a good A_s plays H if he receives a signal that B is weak, and L if he receives a signal that B is tough.

Next, the proof shows that a bad A_s always mixes between L and H . Mixing occurs because the beliefs $\hat{\theta}(e|\mathbf{a})$ vary with the agent's strategy: the more often a bad agent plays tough, the less likely it is (from the seller's perspective) that A_s is good if S observes H , and the more likely it is that he is good if S observes L . Hence $\hat{\theta}(H)$ and $\hat{\theta}(N)$ are decreasing in a_s^b , and $\hat{\theta}(L)$ is increasing in a_s^b . Thus, if for some a_s^b a bad agent prefers H to L because $p\hat{\theta}(H|\mathbf{a}) + (1-p)\hat{\theta}(N|\mathbf{a}) > \hat{\theta}(L|\mathbf{a})$, an increase in a_s^b also reduces the difference between the expected payoffs of playing H or L . The proof shows that an equilibrium is always reached at an interior value of a_s^b where the expected payoffs of H and L are equal.

The remainder of the proof shows that the equilibrium described here is unique. In particular, the requirement that S 's expected payoff be increasing in A_s 's skill rules out all "babbling" equilibria as well as a "perverse" equilibrium in which $\hat{\theta}(N) > \hat{\theta}(L) > \hat{\theta}(H)$ [cf. also the discussion in Scharfstein and Stein (1990)].

A feature of the equilibrium that turns out to be central for our main result is that with open or closed doors, a bad agent plays H less often than does a good agent. Since the expected reputation from playing L does not depend on the agent's type, it must be less than or equal to θ , for otherwise each type of agent could secure an expected reputation greater than θ by playing L , which is not possible. Moreover, since we do not consider pooling equilibria, the expected reputation from playing L must be strictly less than θ . This in turn means that a bad agent must play L at least as often as a good agent, that is $1 - a_s^b > (1-p)q + p(1-q)$ or $a_s^b < pq + (1-p)(1-q)$, as stated in the proposition.

Because of Proposition 1 we have $a^w = 1$ and $a^t = 0$ in equilibrium and hence can further simplify the description of an equilibrium strategy by characterizing only the mixing probability a_s^b , or simply a_s , of a bad agent. Abusing the notation, let $\mathbf{a} = (a_s, a_B)$.

3.1. Optimality of Delegating Bargaining

The equilibrium described in Proposition 3 exists for any $q > 1/2$. If q is relatively small, however, even a good agent is barely better informed about B than S herself is. The seller may then prefer not to delegate bargaining at all, or, if delegation is a practical necessity, may prefer to instruct her agent always to demand a low price or a high price instead of giving him discretion.

Specifically, if S bargained by herself or instructed her agent to demand a certain price, she would choose a high price if

$$p(H - \underline{s}) \geq L - \underline{s}, \quad (6)$$

and otherwise would choose a low price. Thus delegating bargaining to an agent is optimal only if the expected payoff from doing so exceeds the expected payoff from choosing a particular price. Comparing these expected payoffs establishes that delegation is optimal if good agents are sufficiently numerous and well informed:

Proposition 2. Let $\pi_H = H - \underline{s}$ and $\pi_L = L - \underline{s}$. If Equation (6) holds, then delegation is optimal if

$$q \geq \frac{p\pi_H - \pi_L + \theta(1-p)\pi_L}{[p(\pi_H - \pi_L) + (1-p)\pi_L]\theta}.$$

If Equation (6) does not hold, then delegation is optimal if

$$q \geq \frac{(1-p)\theta\pi_L}{2(1-p)\theta\pi_L - (\pi_L - p\pi_H)}.$$

Proposition 2 states sufficient conditions for the optimality of delegation. They are phrased as lower bounds on q , but whether they hold also depends on the fraction of good agents, θ . For example, even if good agents have near-perfect information, delegation will not be optimal if θ is so small that the benefit of letting a good agent make optimal decisions is outweighed by the (from the principal's ex ante perspective) suboptimal behavior of a bad agent according to Proposition 1.

4. Open- versus Closed-Door Bargaining

The main result of this article is

Proposition 3. For any given a_B , a bad agent of a weak seller plays H with higher probability with open doors than with closed doors. As a result, the probability of inefficient disagreement is higher with open doors.

To understand this result, observe that trade at any price is more informative with open than with closed doors because S can see who makes an offer, and because A_s 's price bids reveal information about his type, but his responses to the buyer do not. Thus, with open doors, trade at H is better news about A_s than with closed doors; that is, $\hat{\theta}^o(H) > \hat{\theta}^c(H)$, where the superscripts stand for open and closed doors. Formally, this can be seen by comparing the values of Equation (3) for $i = 0$ for open doors and $i = 1$ for closed doors. Trade at L is worse news about A_s than with closed doors; that is, $\hat{\theta}^o(L) < \hat{\theta}^c(L)$, [cf. Equation (4)]. On the other hand, disagreement leads to the same reputation for open and closed doors. If S is weak, she knows that A_s would accept any offer from B . Disagreement therefore must mean that A_s made an offer, even if S cannot actually observe this; hence $\hat{\theta}^o(N) = \hat{\theta}^c(N)$ [cf. Equation (5)].

These effects combined mean that for a bad agent, bidding H is overall more attractive with open than with closed doors, while bidding L is less attractive. Therefore the a_s at which a bad agent is indifferent between playing H and L must be higher with open than with closed doors. Since this means that trade ends in inefficient disagreement with a higher probability, open-door bargaining is less efficient than closed-door bargaining.

Though derived within very a simple model, our result is likely to be quite general. In particular, it does not depend on the equality of open- and closed-door posteriors when there is no trade. Even though in more general settings this property will no longer hold, there will be an asymmetry between early trade and delay or disagreement in how S 's beliefs differ across the two scenarios; and it is this asymmetry that shifts the agent's behavior in an unambiguous direction. Intuitively, in any kind of multistage bargaining game with (possibly randomly) alternating offers, one would expect the seller agent's acceptance strategy to be relatively more closely pinned to his own principal's reservation price (which the seller knows) and his offering strategy to be more closely pinned to the buyer's reservation price (which the agent wishes to signal he knows).

Thus, from the seller's point of view, with open doors her agent's own offers are more informative about his skill than his responses to the buyer's offers. With closed doors, this distinction cannot be made; and while trade at a high price is good news and trade at a low price bad news about the agent, the seller's beliefs are more compressed toward the prior than the corresponding open-door beliefs. On the other hand, if either the agent's bid is rejected by the buyer or the buyer's bid is rejected by the agent, then bargaining moves on to the next stage or terminates. Irrespective of whose offer was rejected, the seller learns that her agent believes the buyer is weak, and hence updates her belief about the agent in similar ways. Since the seller's closed-door beliefs are, in a sense, averages of her open-door beliefs, depending on which player made an offer, the agent's expected reputation following disagreement or delay will be similar for open- and closed-door bargaining. Both effects combined, that is, the compression of closed-door beliefs when early trade occurs, and the relative lack of compression when delay or disagreement occur, induce a bad agent to bid less aggressively with closed doors, as explained for the one-shot game.

An even more general and simpler intuition that does not depend on the preceding conjectures is the following: when the agent knows his type, a good agent attempts to differentiate himself from a bad agent by following his information, whereas (under either information structure) a bad agent partially imitates a good one by mixing, in spite of his lack of information about the buyer. Now, with open doors, the principal learns more about the agent for any given strategy of a bad agent. Since, as discussed above, a bad agent chooses a high price less often than a good agent, more information increases his expected payoff from choosing a high price and decreases his payoff from choosing a low price. A new equilibrium is thus reached where a bad agent chooses a high price more often than with closed doors. In other words, open-door bargaining induces a bad agent to imitate a good agent to a greater extent.

The last step in solving the game is to determine S 's choice of information structure. We determined above that bargaining on her own, a weak (and uninformed) S would choose a high price if Equation (6) holds. Given Proposition 3, and given that a good agent's strategy does not depend on the information structure, it follows that S will choose open doors if and only if she prefers a bad (uninformed) agent to be tough, the condition for which is again Equation (6). Without further proof, we thus have

Proposition 4. At stage 1 of the game, S chooses open-door bargaining if and only if Equation (6) holds.

Thus for a given behavior of the buyer, open-door bargaining, that is, monitoring her agent, may be unilaterally advantageous for the seller even though it always leads to a less efficient bargaining outcome. What happens when the buyer is a strategic player as well is explored in Section 5.

4.1. Information About the Agents

The seller may want to choose open-door bargaining simply to obtain more accurate information about A_s . It is intuitively clear that with open-door bargaining, S 's information about A_s is more precise for any given bargaining strategies because the set of observable events is strictly finer. In equilibrium, however, open-door bargaining is not necessarily more informative: the bad agent's more aggressive bargaining behavior counterbalances the gain in information, and by example it can be shown that the net outcome is ambiguous. For details, see Fingleton and Raith (2000).

4.2. Reputation as Viewed by Outsiders

The agent might want to maximize his reputation not as perceived by S , but as perceived by an external market. It turns out that in this case, Proposition 3 no longer holds. In our conclusion (Section 6), we discuss the practical implications of this distinction.

Recall that if trade does not occur, a seller with low reservation price knows that A_s must have (unsuccessfully) demanded a high price, for trade would have occurred had he demanded a low price or if the buyer had offered any price. An outside observer who does not know S 's reservation price, however, cannot reach this conclusion: with closed doors, disagreement can occur either if A_s 's bid is rejected by a tough buyer, or if B 's bid is rejected by a tough seller. It can be shown that when the posteriors $\hat{\theta}(e)$ are evaluated from the perspective of an outside observer, a bad agent's response to a change in the information structure no longer has a predictable direction.

We argued above that Proposition 3 is general in that it holds whenever agents choose between a safe and a risky option. What it means to choose a "safe" or a "risky" option, however, depends on the seller's reservation price, which is why the seller's private information matters: a seller with a low reservation price knows that demanding a low price is a safe option, whereas demanding a high price is a risky option. In contrast, even with open doors, an outsider cannot infer from observing a high price bid that the seller's

agent is pursuing a risky strategy; the agent might have had no choice because of the seller's high reservation price. And with closed doors, an outsider's inferences are even more limited.

4.3. Agents Who do not Know Their Type

Proposition 3 is the result of a signaling game between a good and a bad agent. Suppose, instead, that A_s does not know his type,¹³ but only knows that with probability θ he is good, and otherwise bad. If he is bad, he receives a signal that is correct with probability 1/2, that is, it is uninformative. Then, in any equilibrium in which S 's payoff is increasing in A_s 's skill, we must have $\hat{\theta}(H) \geq \hat{\theta}(L) \geq \hat{\theta}(N)$, with at least one inequality strict. Since A_s is good with positive probability, his signal is informative in ex ante terms and his probability of playing tough must be weakly increasing in the signal he gets about B 's reservation price.

Several types of equilibria are possible in this game, depending on the parameters of the model. If p is large or small, A_s may never use his signal and instead always plays H or L . If he does use his signal, however, then he will play H if and only if he receives a signal that B is weak. Since the probability of this event does not depend on the information structure, there is no difference in behavior between open- and closed-door bargaining.¹⁴

5. Bargaining with a Strategic Buyer

We have so far focused on A_s 's best response to a given strategy used by B . We now turn to the full bargaining game in which the buyer (or her agent) also behaves strategically.

5.1. Direct Bargaining with the Buyer

Suppose A_s negotiates directly with B , but now assume that B acts strategically. In our one-shot model, the players' strategies are unrelated because each side accepts any offer made by the other as long as it is compatible with the reservation price. Hence, whether bargaining takes place with open or closed doors has no effect on B 's strategy, and the results of Section 4 apply: with open-door bargaining, a bad A_s bargains more aggressively, and the expected bargaining outcome is less efficient. The seller nevertheless gains from monitoring if she wants a bad agent to be tough, that is, if Equation (6) holds.

More generally, however, if one party makes a stronger demand, the other party's best response usually is to concede more easily, to avoid delay or disagreement. This is most obvious in a simple Nash demand game, but also holds for bargaining under incomplete information. Open-door bargaining may

13. This assumption is made in Holmström (1999), Scharfstein and Stein (1990), and the basic models of Ottaviani and Sørensen (2000) and Prat (2003).

14. The formal conditions that lead to one or the other type of equilibrium are difficult to interpret. It is possible that switching from closed to open doors may lead to a shift from a pooling to a separating equilibrium, or vice versa. The direction of change in the agent's behavior also seems ambiguous.

therefore be a way for the seller to credibly commit to a more aggressive bargaining position and thereby gain an advantage over the buyer, as pointed out by Schelling (1960:29–30):

A potent means of commitment . . . is the pledge of one's reputation. If national representatives can arrange to be charged with appeasement for every small concession, they place concession visibly beyond their own reach . . . But to commit in this fashion publicity is required. Both the initial offer and the final outcome would have to be known; and if secrecy surrounds either point . . . the device is unavailable.

Our model is the first to formally link the idea of commitment through delegation to an agent's reputation and the role of openness, thereby avoiding the usual problem of credibility. It is obvious that the seller may find it advantageous to commit to a certain (aggressive) bargaining position if the buyer's best response is to be more appeasing. The problem, however, is that a principal's instructions to her agent to be a tough bargainer fail to be credible if principal and agent have an incentive to secretly renegotiate. In general, commitment through a contract is credible only if the agent's objectives are genuinely different from the principal's, that is, in the presence of a genuine agency problem (see Katz, 1991; Caillaud and Rey, 1994).¹⁵

Most models of delegated bargaining assume that an agency problem exists because of a difference between the principal's and the agent's discount rate or risk aversion, that is, a preference parameter that affects the outcome of bargaining [cf. Jones (1989), Burtraw (1993), and Segendorff (1998)]. No such differences are assumed here; bargaining itself is a cheap-talk game. What gives rise to an agency problem is that the principal cares about the bargaining outcome, but the agent only about his reputation. Commitment is achieved not by choosing an agent with particular preferences, but by inducing more aggressive bargaining as equilibrium behavior as a result of monitoring.

5.2. Bargaining Between Two Agents

Often, both parties engaged in bargaining are represented by agents. To study this situation, assume now that A_s bargains not directly with B , but with an agent A_b hired by B , about whom we make the same assumptions as we did about A_s .¹⁶ We can then show:

Proposition 5. In the bargaining game between A_s and A_b , there exists a unique and symmetric equilibrium for both open- and closed-door bargaining. Bad types of both agents play tough with probability a^o and a^c ,

15. Conditions under which delegation can be effective even when contracts are observable are studied in Fershtman and Kalai (1997) and Kockesen and Ok (2002). Bester and Sákovics (2001) study the case of costly renegotiation.

16. In our one-shot model, it is irrelevant whether the agents know each other's type. In a more general game, whether they do would matter for the calculations, but probably not for the validity of Proposition 3, in light of the general intuition given above.

respectively, where $a^o, a^c \in (0, 1)$. Finally, agents bargain more aggressively with open doors, that is, $a^o > a^c$, leading to a less efficient bargaining outcome, than with closed doors.

Given Proposition 1, the existence of a symmetric interior equilibrium in the game between two agents is straightforward. Uniqueness is established by showing that the “reaction function” $a_i^b(a_j^b)$ for the bad type of each side’s agent has a slope of less than one at any symmetric point $a_s^b = a_b^b = a$. Finally, the result that $a^o > a^c$ follows from Proposition 3.

It is easy to relate the parameter a_B used earlier to A_b ’s equilibrium strategy in a game with two-sided delegation. In equilibrium, a good A_b uses his information efficiently, whereas a bad A_b mixes; let a_b^b denote that probability that a bad A_b plays L . If S is weak, a good A_b (who receives a signal about S) plays L with probability q . Thus, from A_s ’s perspective, the probability that the agent of a weak buyer plays L equals $a_B = \theta q + (1 - \theta)a_b^b$.

Proposition 5 formalizes the common notion that open-door bargaining induces agents to engage in posturing and therefore leads to an inefficient bargaining process. Anticipating this outcome, principals therefore have an incentive to let bargaining proceed behind closed doors; that is, to commit themselves not to monitor their agents.¹⁷

As in the single-agent case, open-door bargaining may or may not provide more information about agents in equilibrium. If open-door bargaining is more informative, the principals face a trade-off between the efficiency of the current bargaining situation and the efficiency of future bargaining situations. If, however, the advantages of transparency in open-door bargaining are entirely undone by the agents’ more aggressive bargaining, the principals unambiguously prefer closed doors.

5.3. Two Agents, Unilateral Monitoring

Schelling’s quote above suggests that open-door bargaining may arise as the outcome of a prisoners’ dilemma in which each principal individually has an incentive to monitor the agents. We have argued in Section 5.1 that by monitoring the agents, the seller may gain an advantage due to his agent’s more aggressive behavior. What if the buyer, too, is represented by an agent concerned with his reputation? To examine this question, consider the reverse situation: imagine that B can decide to sit in the room in which the agents negotiate, while S is outside and learns only the eventual outcome. Does B have an incentive to monitor the agents if S does not monitor them? And does S have an incentive to monitor the agents as well if B already monitors?¹⁸

17. The Schelling quotation above continues: “If both parties fear the potentialities for stalemate in the simultaneous use of this tactic, they may try to enforce an agreement on secrecy.”

18. We continue to assume here that S and B choose whether to monitor the agents before they learn their reservation prices, that is, before a particular bargaining situation arises. We thus do not consider situations in which a principal’s monitoring decision is a signal of her reservation price.

If B monitors the agents, the situation for her is as with open doors. In this case, A_b 's equilibrium strategy is a^o according to Proposition 5 and does not depend on A_s 's strategy. Since $a^o > a^c$, a bad A_b bargains more aggressively than behind closed doors if he is monitored by B . In the absence of a change in A_s 's behavior, monitoring is profitable for B if she prefers a bad agent to be tough, that is, if

$$p(\bar{b} - L) > \bar{b} - H, \quad (7)$$

which is the counterpart of Equation (6) for the buyer.

While in general a seller negotiating with A_b directly might be more appealing if A_b is more aggressive, the optimal response of A_s is of ambiguous direction. To see why, notice that S faces the same situation as with closed doors; hence A_s 's best response is given by $a_s^b(a_B)$ according to Proposition 1. How a_s changes with a_B can be seen by examining the posterior beliefs $\theta(e)$ as given by Equation (A1), with $i = 1$.

Inspection shows that both $\hat{\theta}(H)$ and $\hat{\theta}(L)$ are increasing in a_B , whereas $\hat{\theta}(N)$ does not depend on a_B . Intuitively, $\hat{\theta}(H)$ increases with a_B because if A_b is more likely to bid L , then trade at H is more likely to be brought about by A_s . On the other hand, a higher a_B also raises the chances that trade at L is brought about by A_b . This increases the "dilution" of $\hat{\theta}(L)$ due to offers made by A_b and moves it closer to the prior θ , which means that $\hat{\theta}(L)$ must increase since $\hat{\theta}(L) < \theta$. These two effects move a_s in opposite direction. It can be shown that in our model, the first effect dominates for any a_B at or above the closed-door equilibrium level, implying that a_s is increasing in a_B in the relevant range. It is unclear, however, how general this result is.

Overall, three different situations can arise: first, if Equation (7) does not hold, then B prefers a bad agent to be appeasing, and therefore prefers not to monitor the agents. Second, if Equation (7) holds, B prefers A_b to be tough other things being equal, but if A_s responds with more aggressive behavior, monitoring might still not be profitable for B . Finally, if monitoring is profitable for B and if Equation (6) holds (in particular, if the bargaining game is symmetric), then S would want to reciprocate by monitoring as well. The outcome would be open-door bargaining, and both parties would be unambiguously worse off than without any monitoring. That is, the principals would face a prisoners' dilemma, as Schelling (1960) envisioned.

6. Conclusion

We have argued that when parties engaged in bargaining are represented by agents with career concerns, the agents' attempts to promote their reputation lead to bargaining behavior that is neither profit- nor surplus-maximizing. This result already holds when agents bargain behind closed doors, but open-door bargaining exacerbates the tendency of agents to engage in "posturing." We have also argued that this result holds generally in settings in which agents signal their knowledge about the state of the world through their choice between risky and safe strategies.

To isolate the consequences of reputation concerns for bargaining, we have abstracted from other forms of moral hazard. For example, achieving an efficient agreement may also require effort on the part of the agents, and monitoring may be necessary to induce this effort. Principals must then weigh the agents' reputation concerns against their incentive to shirk and against the principals' interest in learning about their agents. Other things being equal, principals are more likely to prefer closed-door bargaining:

(1) If an efficient agreement in the current bargaining situation is very important compared with the evaluation of the agents' performance. Hence, diplomacy in international crises is more likely to be surrounded by secrecy than regular discourse in national politics.

(2) If an agent's reputation strongly depends on his current performance as a negotiator. For this reason, managers, who are evaluated on many criteria other than the outcome of labor negotiations, are less likely to engage in posturing than are union leaders. Indeed, studies have shown that negotiators with a strong reputation are less likely to be sanctioned by their constituents (Pruitt, 1981:43–44).

(3) If agents care more about how they are perceived by their own principals than by an external market, since in this case posturing is most likely to occur. For example, politicians, diplomats, and union leaders are concerned with their reputations within their own organizations (their party, country, union) and not different ones, and are therefore likely to engage in posturing when observed by their constituents. Accordingly, closed-door negotiations in politics, diplomacy, and labor negotiations are very common. In contrast, if an agent's market reputation rests more on what the market can observe (for instance, attorneys' observed performance in court), posturing is less likely to occur, and thus closed-door trials in court seem difficult to justify on grounds of efficiency.

(4) If an agent's ability and effort are complementary. In this (realistic) case, high-skill agents are those with a low cost of effort, and then the situation is essentially as described by our model. If, on the other hand, agents can easily compensate for a lack of skill by working harder, monitoring may be desirable.

(5) If it is relatively easy to determine whether a bargaining outcome is efficient. In our model, trade at any price is efficient if it occurs, while disagreement may or may not be. In contrast, if principals cannot assess the efficiency of an agreement because the bargaining situation is too complex (sometimes referred to as "integrative bargaining" in the negotiation literature), they also cannot determine whether the agents' efforts were adequate (or whether they colluded). Monitoring may then be desirable.

(6) If outcomes are sufficiently informative of an agent's bargaining behavior. It may be easy for a principal to hold a single agent accountable for the results he achieves. If the principal is represented by a team of delegates, however, free-rider problems arise that may make monitoring necessary. For example, as Kroszner and Stratmann (2000) point out, the degree of

transparency in political decision making affects the ability of interest groups to monitor politicians' voting behavior and hence the politicians' ability to raise money from interest groups.

For better or worse, those regularly involved in bargaining often attempt to thwart efforts by their principals to increase transparency. For instance, a common argument against televising parliamentary debates is the alleged consequence that "all real debate will be carried out in the corridors" (von Sydow, 1995). In wage bargaining, delegates from both sides whose moves are monitored often engage in tacit communication, that is, use "language which has different meanings to different audiences" in order to communicate information that "cannot be stated explicitly without risk of severe censure from the respective constituent groups" (Walton and McKersie, 1965:336–37). Such responses are sometimes seen as an attempt by agents to escape accountability. Our theory suggests, though, that they may also be aimed at reducing the pressure to engage in posturing.

Elements of our theory appear to generalize to situations other than formal bargaining settings and situations where people do not officially act as agents of others. Debates and discussions often seem more productive when held in private than in public settings. The reason is that discussions without observers tend to be concerned with finding a solution or common ground, whereas the presence of observers introduces a second, distracting, motive: the desire to win an argument.

Appendix: Proofs

Proof of Proposition 1. (1) It will be convenient to consider the cases of open- and closed-door bargaining simultaneously where possible. Let the indicator variable i equal zero for open-door and one for closed-door bargaining. From S 's point of view, the probability that a good A_s trades at H is $p[i(1 - a_B) + qa_s^w + (1 - q)a_s^t]/2$, where the factor p accounts for the fact that H can occur only if B is weak. The distinction between $i = 0$ and $i = 1$ accounts for the difference between the event H with open doors, where S observes that A_s made an offer, and H with closed doors, where S cannot observe who made a bid H (a weak B does so with probability $1 - a_B$). Similarly the probability that a bad A_s trades at H is $p[i(1 - a_B) + a_s^b]/2$. Applying Bayes' rule, the probability that A_s is good when H occurs is

$$\hat{\theta}(H) = \frac{[i(1 - a_B) + qa_s^w + (1 - q)a_s^t]\theta}{[i(1 - a_B) + qa_s^w + (1 - q)a_s^t]\theta + [i(1 - a_B) + a_s^b](1 - \theta)},$$

which leads to Equation (3). Event N can occur only if B is tough, and only if A_s bids H . A bad A_s will do so with probability a_s^b , a good A_s with probability $qa_s^t + (1 - q)a_s^w$. Therefore S 's beliefs are given by Equation (5). For event L we need to distinguish between a tough or a weak B . If B is tough, she always

bids L , while a good A_s bids L with probability $q(1 - a_s^w) + (1 - q)(1 - a_s^t)$. If B is weak, she bids L with probability a_B , and a good A_s with probability $q(1 - a_s^t) + (1 - q)(1 - a_s^w)$. A bad A_s bids L with probability $1 - a_s^b$, irrespective of B 's type. This leads to Equation (4).

(2) Next, we show that if $\hat{\theta}(H) > \theta > \hat{\theta}(L) > \hat{\theta}(N)$, then there exists a unique equilibrium in which a good player bids according to his signal, as described in the proposition. Since the payoff of Equation (1) is linear in p^k , at most one of the three types w, b, t can be indifferent between L and H . This cannot be either of the good types of A_s : with open doors, if type w mixes, that is, $a_s^w < 1$, then in equilibrium a_s^b and a_s^t must both equal zero. But then both $\hat{\theta}(H)$ and $\hat{\theta}(N)$ would equal one, in which case the agent would certainly prefer H over L , contradicting the assumption. A similar argument applies for a type t agent.

For closed-door bargaining, some calculations are required, but the argument is the same: if $a_s^w < 1$ and $a_s^b = a_s^t = 0$, then the resulting $\hat{\theta}(H)$ is increasing in a_s^w , $\hat{\theta}(H)$ equals one, and $\hat{\theta}(L)$ is decreasing in a_s^w . At $a_s^w = 0$, $\hat{\theta}(L)$ equals the prior θ , whereas the expected payoff from playing H equals $[1 - q - p(1 - q(1 + \theta))]/[1 - p - q + 2pq]$, which exceeds θ . It follows that for any a_s^w used in posterior beliefs, a type w agent would prefer to play H , contradicting the assumption. Finally, if $a_s^t > 0$ and therefore $a_s^b = a_s^w = 1$, then the resulting $\hat{\theta}(H)$ and $\hat{\theta}(N)$ are increasing in a_s^t and $\hat{\theta}(H)$ is decreasing in a_s^t . At $a_s^t = 1$, the payoffs of playing H or L both equal θ , which means that in any separating equilibrium, a type t agent strictly prefers to play L , contradicting the assumption.

(3) As a result of step 2, we can from now on set $a_s^w = 1$ and $a_s^t = 0$. The simplified expressions for S 's beliefs are

$$\begin{aligned} \hat{\theta}(H) &= \frac{[(1 - a_B)i + q]\theta}{i(1 - a_B) + q\theta + a_s^b(1 - \theta)}, \\ \hat{\theta}(N) &= \frac{(1 - q)\theta}{(1 - q)\theta + a_s^b(1 - \theta)}, \\ \hat{\theta}(L) &= \frac{[(1 - p + pa_B)i + (p + q - 2pq)]\theta}{(1 - p + pa_B)i + (p + q - 2pq)\theta + (1 - a_s^b)(1 - \theta)}. \end{aligned} \tag{A1}$$

An equilibrium in which a bad A_s mixes exists if for some a_s^b the payoff from playing H , $p\hat{\theta}(H) + (1 - p)\hat{\theta}(N)$, equals the payoff $\hat{\theta}(L)$ if A_s plays L . By inspection, $\hat{\theta}(H)$ and $\hat{\theta}(N)$ are decreasing in a_s^b , whereas $\hat{\theta}(L)$ is increasing in a_s^b . It follows that the difference

$$\Delta\pi(a_s^b) = p\hat{\theta}(H) + (1 - p)\hat{\theta}(N) - \hat{\theta}(L) \tag{A2}$$

is strictly decreasing in a_s^b , and can therefore have at most one root in $(0, 1)$. For open-door bargaining we have

$$\Delta\pi(0) = \frac{1 - \theta}{1 - \theta(1 - p - q + 2pq)} > 0$$

and

$$\Delta\pi(1) = -\frac{(1 - \theta)[1 - \theta(1 - p - q + 2pq)]}{(1 - q\theta)(1 - \theta + q\theta)} < 0.$$

With closed doors, for $a_s^b = 0$ we have $\hat{\theta}(N) = 1 > \hat{\theta}(H)$, and

$$\hat{\theta}(H) - \hat{\theta}(L) = \frac{\{1 + q - [p(1 - a_B) + a_B](1 - q)\}\theta(1 - \theta)}{(1 - a_B + q\theta)[2 - p + pa_B - \theta(1 - p - q + 2pq)]} > 0.$$

It follows that $\Delta\pi(0) > 0$. At $a_s^b = 1$ we have

$$\hat{\theta}(H) - \hat{\theta}(N) = \frac{(3q - 1 - a_Bq)\theta(1 - \theta)}{(2 - a_B - (1 - q)\theta)(1 - q\theta)} > 0$$

and

$$\hat{\theta}(L) - \hat{\theta}(H) = \frac{\{1 - q[a_B + p(1 - a_B)] + p + q - 2pq\}\theta(1 - \theta)}{[2 - a_B - \theta(1 - q)][1 - p + pa_B + \theta(p + q - 2pq)]} > 0,$$

and therefore $\Delta\pi(1) < 0$. Since with both open and closed doors, $\Delta\pi$ is positive at $a_s^b = 0$ and negative at $a_s^b = 1$, $\Delta\pi$ is unique in the interior of the unit interval, which proves the existence and uniqueness of the equilibrium.

(4) Next we show that in equilibrium, $a_s^b < 1 - p - q + pq$. Using Equation (A1) and substituting $1 - p - q + pq$ for a_s^b , $\Delta\pi$ reduces to

$$\frac{p(1 - p)\theta(1 - \theta)(2q - 1)[\theta(2q - 1) + (1 - a_B)i]}{[1 - q + p(1 - \theta)(2q - 1)][1 - q + (1 - a_B)i + (2q - 1)(\theta + p - p\theta)]} < 0.$$

Since $\Delta\pi$ is decreasing in a_s^b , the equilibrium value of a_s^b must be smaller than $1 - p - q + pq$ for both open and closed doors.

(5) Given the agent's strategy determined above, the resulting beliefs satisfy $\hat{\theta}(H) > \theta > \hat{\theta}(L) > \hat{\theta}(N)$: first, we have

$$\hat{\theta}(H) - \theta = \frac{(q - a_s^b)\theta(1 - \theta)}{(1 - a_B)i + q\theta + a_s^b(1 - \theta)},$$

which is positive because $a_s^b < 1 - p - q + pq < q$, where the second inequality follows from $q > 1/2$. Next, we have

$$\theta - \hat{\theta}(L) = \frac{(1 - p - q + 2pq - a_s^b)\theta(1 - \theta)}{(1 - p + pa_B)i + (p + q - 2pq)\theta + (1 - a_s^b)(1 - \theta)} > 0.$$

These steps imply that $\hat{\theta}(H) > \hat{\theta}(L)$ for any a_s^b . Since a_s^b satisfies $\Delta\pi = 0$, it follows that in equilibrium, $\hat{\theta}(L)$ must exceed $\hat{\theta}(N)$ for both open- and closed-door bargaining.

(6) There exist two other kinds of equilibrium, both of which violate the requirement that on average a good agent negotiates a higher payoff for S than a bad agent. First, since bids are cheap talk, any pooling strategy with $a_s^t = a_s^b = a_s^w = a$ can be part of an equilibrium, supported by beliefs that treat each event as uninformative, that is, $\theta(e) = \theta$ for all e that occur with positive probability, and $\hat{\theta}(e) \leq \theta$ for events that occur with zero probability. Second, there exists a “perverse” equilibrium in which $a_s^t = 1$, $a_s^w = 0$, and $a_s^b \in [0, 1)$; and $\hat{\theta}(N) > \theta > \hat{\theta}(L) > \hat{\theta}(H)$, which is also a common feature of cheap-talk games [cf. Scharfstein and Stein (1990)].

(7) There exist no separating equilibria other than those of the proposition and its perverse counterpart because the beliefs must be ranked $\hat{\theta}(H) > \hat{\theta}(L) > \hat{\theta}(N)$ or the other way round. No weak inequalities and no other orders are possible. Suppose, for example, that $\hat{\theta}(H) = \hat{\theta}(L) > \hat{\theta}(N)$, all agents would strictly prefer L (because even a w type may by mistake end up without trade); that is, there would be pooling. And if, for example, the beliefs were $\hat{\theta}(H) > \hat{\theta}(N) > \hat{\theta}(L)$, then playing H would be a dominant strategy for all players; that is, again there would be pooling. That is, since playing L leads to the event L for all types of agents, L can in equilibrium never be the best or the worst outcome. ■

Proof of Proposition 2. The seller’s expected payoff from delegating bargaining to an agent of unobservable ability is given by

$$\pi_D = \theta \{ p[q\pi_H + (1 - q)\pi_L + (1 - p)q\pi_L] + (1 - \theta) \times \{ p[a_s\pi_H + (1 - a_s)\pi_L] + (1 - p)(1 - a_s)/\pi_L \}.$$

If Equation (6) holds, then delegation is optimal if $\Delta_H := \pi_D - p\pi_H > 0$. Since Δ_H is increasing in q , the condition $\Delta_H > 0$ can be rephrased as

$$q \geq \frac{(1 - a_s + \theta a_s)(p\pi_H - \pi_L) + (1 - p)\theta\pi_L}{[p(\pi_H - \pi_L) + (1 - p)\pi_L]\theta}. \tag{A3}$$

$$q \geq \frac{a_s(1 - \theta)(\pi_L - p\pi_H) + (1 - p)\theta\pi_L}{[p(\pi_H - \pi_L) + (1 - p)\pi_L]\theta}. \tag{A4}$$

The lower bound in Equation (A3) is decreasing in a_s and so, a sufficient condition for Equation (A3) is obtained by substituting 0 for a_s . Rearranging for q then leads to the lower bound stated in the proposition.

If Equation (6) does not hold, then delegation is optimal if $\Delta_L := \pi_D - \pi_L > 0$. Since Δ_L is increasing in q , this condition can be rephrased as The lower bound in Equation (A4) is increasing in a_s . The arguments used above again apply, and so a sufficient condition for Equation (A4) is obtained by substituting q for a_s in Equation (A4). Rearranging for q leads to the stated lower bound.

Proof of Proposition 3. Recall the notation introduced after Proposition 1 whereby a_s is a shorthand for the strategy a_s^b of a bad agent, whereas in equilibrium a_s^o and a_s^c are one and zero, respectively. Let the superscripts “o” and “c” stand for open- and closed-door bargaining, respectively. From Proposition 1, we know that both a_s^o and a_s^c and hence any $a_s \in [\min\{a_s^o, a_s^c\}, \max\{a_s^o, a_s^c\}]$ is less than $1 - p - q + 2pq$. Using Equation (A1), the derivative of $\hat{\theta}(H)$ with respect to i is

$$-\frac{(1 - a_B)(q - a_s^b)\theta(1 - \theta)}{[(1 - a_B)i + q\theta + a_s^b(1 - \theta)]^2},$$

which is negative because $a_s < 1 - p - q + 2pq$ implies $a_s < q$. The derivative of $\hat{\theta}(L)$ with respect to i is

$$\frac{(1 - p + pa_B)(1 - p - q + 2pq - a_s^b)\theta(1 - \theta)}{[(1 - p + pa_B)i + (p + q - 2pq)\theta + (1 - a_s^b)(1 - \theta)]^2} > 0.$$

Finally, by inspection, $\hat{\theta}(N)$ does not depend on i . It follows that $\Delta\pi$ according to Equation (A2) must be decreasing in i in the relevant range of a_s . Since $\Delta\pi$ is also strictly decreasing in a_s , the derivative $da_s/di(\Delta\pi = 0)$ is negative, which means $a^o > a^c$.

Proof of Proposition 5. First, Proposition 1 holds analogously for the buyer’s side, so in equilibrium, a good A_b uses his information efficiently, and a bad A_b plays tough (L) with some probability $a_b^b \in (0, 1)$. In particular, if S is weak, then a good A_b plays L with probability q . Thus the probability from S ’s point of view that a weak buyer plays L is given by $a_B = \theta q + (1 - \theta)a_b^b$. Since the resulting functions $a_s^b(a_b^b)$ and $a_b^b(a_s^b)$ are continuous and lie in the interior of the unit interval, a point of intersection must exist. Moreover, since the model is symmetric, any intersection must be at a symmetric point $a_s^b = a_b^b = a$.

With open doors, each agent’s strategy is independent of the other’s, so the equilibrium must be unique. For closed doors, a sufficient condition for uniqueness is that at any point of intersection, the slope $da_s^b(a_b^b)/da_b^b$ is less than one. Using $a_B = \theta q + (1 - \theta)a_b^b$ and $i = 1$ in Equation (A1) and then Equation (A2), $a_s^b(a_b^b)$ is the solution to

$$\Delta(a_s^b, a_b^b) = \frac{p[1 + q - \theta - (1 - \theta)a_b^b]}{1 - \theta - (1 - \theta)a_b^b + q\theta + a_s^b(1 - \theta)} + \frac{(1 - p)(1 - q)}{(1 - q)\theta + a_s^b(1 - \theta)} - \frac{1 - p(1 - \theta)(1 - a_b^b) + p + q - 2pq}{1 - p(1 - \theta)(1 - a_b^b) + (p + q - 2pq)\theta + (1 - a_s^b)(1 - \theta)} = 0.$$

Since $da_s^b/da_b^b = -(\partial\Delta/\partial a_b^b)/(\partial\Delta/\partial a_s^b)$, the condition $da_s^b/da_b^b < 1$ is equivalent to $\partial\Delta/\partial a_b^b + \partial\Delta/\partial a_s^b < 0$. To see that this is the case, evaluate $\partial\Delta/\partial a_b^b$ at $a_b^b = a_s^b = a$:

$$p(1-\theta)^2 \times \left\{ \frac{q-a}{[1-(1-q)\theta]^2} - \frac{1-p-q+2pq-a}{[1-p(1-\theta)(1-a)+(p+q-2pq)\theta+(1-a)(1-\theta)]^2} \right\},$$

and evaluate $\partial\Delta/\partial a_s^b$ at $a_b^b = a_s^b$:

$$-(1-\theta) \left\{ \frac{p[1-\theta+q-(1-\theta)a]}{[1-(1-q)\theta]^2} + \frac{(1-p)(1-q)}{[(1-q)\theta+a(1-\theta)]^2} \right. \\ \left. + \frac{1-p(1-\theta)(1-a)+p+q-2pq}{[1-p(1-\theta)(1-a)+(p+q-2pq)\theta+(1-a)(1-\theta)]^2} \right\}.$$

The sum of these is the negative of

$$\frac{p}{1-(1-q)\theta} + \frac{(1-p)(1-q)}{[(1-q)\theta+a(1-\theta)]^2} \\ + \frac{1-pq(1-\theta)+(p+q-2pq)+p^2(1-\theta)(2q-1)}{[1-p(1-\theta)(1-a)+(p+q-2pq)\theta+(1-a)(1-\theta)]^2},$$

in which each term is negative. The steps so far establish the existence and uniqueness of an equilibrium with mixing probabilities of a^o and a^c , respectively. Finally, Proposition 3 implies that a^o must exceed a^c , and since disagreement is more likely to occur if both agents bargain more aggressively, the open-door equilibrium is less efficient. ■

References

Bacon, Francis. 1862 [1625]. *Bacon's Essays and Colours of Good and Evil*. London: MacMillan.

Bester, Helmut, and József Sákovics. 2001. "Delegated Bargaining and Renegotiation," 45 *Journal of Economic Behavior and Organization* 459–73.

Burtraw, Dallas. 1993. "Bargaining With Noisy Delegation," 24 *Rand Journal of Economics* 40–57.

Cai, Hongbin. 2000. "Bargaining on Behalf of a Constituency," 92 *Journal of Economic Theory* 234–73.

Caillaud, Bernard, and Patrick Rey. 1994. "Strategic Aspects of Vertical Delegation," 39 *European Economic Review* 421–31.

Chatterjee, Kalyan, and Larry Samuelson. 1987. "Bargaining with Two-Sided Incomplete Information: An Infinite Horizon Model with Alternating Offers," 54 *Review of Economic Studies* 175–92.

Court of First Instance of the European Communities. 1995. Judgement in Case T-194/94 (October 19).

Cramton, Peter. 1985. "Sequential Bargaining Mechanisms," in Alvin Roth, ed., *Game-Theoretic Models of Bargaining*. Cambridge: Cambridge University Press.

Crawford, Vincent P. 1982. "A Theory of Disagreement in Bargaining," 50 *Econometrica* 607–38.

Fershtman, Chaim, and Ehud Kalai. 1997. "Unobserved Delegation," 38 *International Economic Review* 763–74.

Fingleton, John, and Michael Raith. 2000. "Open Covenants, Privately Arrived At," mimeo.

Fisher, Roger, and William Ury. 1981. *Getting to Yes*, rev. ed. London: Century Press.

Holmström, Bengt. 1999. "Managerial Incentive Problems: A Dynamic Perspective," 66 *Review of Economic Studies* 169–82.

- Jones, Stephen R. G. 1989. "Have Your Lawyer Call My Lawyer—Bilateral Delegation in Bargaining Situations," 11 *Journal of Economic Behavior and Organization* 159–74.
- Katz, Michael. 1991. "Game-Playing Agents: Unobservable Contracts as Precommitments," 22 *RAND Journal of Economics* 307–28.
- King, David C., and Richard J. Zeckhauser. 1999. "Legislators as Negotiators," in Robert H. Mnookin and Lawrence E. Susskind, eds., *Negotiating on Behalf of Others*. Thousand Oaks, CA: Sage Publications.
- Kockesen, Levent, and Efe Ok. 2002. "Strategic Delegation by Unobservable Incentive Contracts," Economics Discussion Paper 0102–26, Columbia University.
- Kroszner, Randall S., and Thomas Stratmann. 2000. "Does Political Ambiguity Pay? Corporate Campaign Contributions and the Rewards to Legislator Reputation," Working Paper 7475, NBER.
- Kurtzberg, Terri, Don Moore, Kathleen Valley, and Max Bazerman. 1999. "Agents in Negotiations: Towards Testable Predictions," in Robert H. Mnookin and Lawrence E. Susskind, eds., *Negotiating on Behalf of Others*. Thousand Oaks, CA: Sage Publications .
- Mnookin, Robert H., and Jonathan R. Cohen. 1999. "Introduction," in Robert H. Mnookin and Lawrence E. Susskind, eds., *Negotiating on Behalf of Others*. Thousand Oaks, CA: Sage Publications.
- Myerson, Roger, and Mark Satterthwaite. 1983. "Efficient Mechanisms for Bilateral Trading," 29 *Journal of Economic Theory* 265–81.
- Ottaviani, Marco, and Peter Sørensen. 2000. "Professional Advice," mimeo, University College London and University of Copenhagen.
- Perry, Motty, and Larry Samuelson. 1994. "Open- versus Closed-Door Negotiations," 25 *RAND Journal of Economics* 348–59.
- Prat, Andrea. 2003. "The Wrong Kind of Transparency," mimeo, London School of Economics.
- Prendergast, Canice, and Lars Stole. 1996. "Impetuous Youngsters and Jaded Oldtimers," 104 *Journal of Political Economy* 1105–34.
- Pruitt, Dean G. 1981. *Negotiation Behavior*. New York: Academic Press.
- Rubin, Jeffrey Z., and Bert R. Brown. 1975. *The Social Psychology of Bargaining and Negotiation*. New York: Academic Press.
- Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Scharfstein, David, and Jeremy Stein. 1990. "Herd Behavior and Investment," 80 *American Economic Review* 465–79.
- Segendorff, Björn. 1998. "Delegation and Threat in Bargaining," 23 *Games and Economic Behavior* 266–83.
- von Sydow, Emily. 1995. "Clamour for Openness," 1 *European Voice* (9–15).
- Walton, Richard E., and Robert B. McKersie. 1965 [1991]. *A Behavioral Theory of Labor Negotiations*. Ithaca, NY: ILR Press.