DISCUSSION PAPER SERIES

No. 2709

OPEN COVENANTS, PRIVATELY ARRIVED AT

John Fingleton and Michael Raith

INDUSTRIAL ORGANIZATION

Centre for Economic Policy Research

www.cepr.org

Available online at: www.cepr.org/pubs/dps/DP2709.asp

ISSN 0265-8003

OPEN COVENANTS, PRIVATELY ARRIVED AT

John Fingleton, Irish Competition Authority
Michael Raith, University of Chicago and CEPR

Discussion Paper No. 2709
February 2001

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 2709

February 2001

ABSTRACT

Open Covenants, Privately Arrived At*

This Paper analyses strategic bargaining between two agents each of whom negotiates on behalf of a principal. The principals face uncertainty about the bargaining skills of their agents as measured by the agents' abilities to assess the opponent's preferences. Agents then have an incentive to promote their reputation as skilled bargainers through their bargaining behaviour. We compare two different scenarios: open-door bargaining, where the principals observe the entire bargaining process, and closed-door bargaining, where they observe only the final outcome. We show that with open doors, the higher visibility of agents' actions induces low-skill agents to negotiate more aggressively than behind closed doors in order to distort their principals' inferences. Since this 'posturing' increases the probability of delay or disagreement, closed-door bargaining is more efficient.

JEL Classification: C78, D70, D82
Keywords: career concerns, delegated bargaining, reputation, strategic bargaining

John Fingleton
Competition Authority
Parnell House
14 Parnell Square
Dublin 1
IRELAND
Tel: (353 1) 804 5400
Fax: (353 1) 804 5401
Email: john.fingleton@attglobal.net

Michael Raith
Graduate School of Business
University of Chicago
1101 E. 58th Street
Chicago, IL 60637
USA
Tel: (1 773) 702 7324
Fax: (1 773) 702 0458
Email: michael.raith@gsb.uchicago.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid= 101021

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid= 139602

*We thank Canice Prendergast for his advice and support, and Judy Chevalier, Guido Friebel, Luis Garicano, Allison Garrett, Paul Heidhues, Randy Kroszner, Matthias Raith, Philip Reny, Lars Stole, and seminar participants at Chicago GSB, the Stockholm School of Economics, and at conferences for many helpful comments. Any errors are our own. We also thank 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 (ECARES) in Brussels where this research was initiated. Raith acknowledges financial support from the University of Chicago Graduate School of Business.

Submitted 27 November 2000

NON-TECHNICAL SUMMARY

In most bargaining situations in business and politics, negotiators bargain not for themselves, but on behalf of constituents, or principals. Two views exist on how closely agents in these situations ought to be monitored by those who appoint them. Some maintain that agents' actions should always be as transparent as possible to the principals represented and that negotiations should thus be conducted with 'open doors'. Others argue that too much monitoring in some way distracts agents and leads to *less* efficient decision-making processes. In their view, 'closed-door' negotiations allow principals to hold their agents accountable for the *results* they negotiate without interfering with the *process* of negotiation.

These two views are reflected in a large variety of institutional arrangements. Labour negotiations often start with a public exchange of demands between managers and union leaders, but are concluded in closed-door meetings. In parliaments, debates may be televised (as in the US) or held behind closed doors (as in the UK until 1989); and voting records may be public (as in Sweden) or secret (as in Germany). Such differences between countries have fuelled the recent debate over the appropriate degree of transparency in EU decision-making. While delegates routinely emphasize the inefficiencies caused by openness, constituents often suspect that such arguments may be self-serving.

From an agency-theoretic perspective, the benefits of monitoring agents are obvious: monitoring facilitates the rewarding of agents for their performance and thus induces agents to work harder. It can also help prevent collusion between agents. Less obvious is why monitoring can have detrimental effects. The purpose of this Paper is to provide an answer to this question, and thus an explanation for why even in democratic institutions closed-door negotiations are prevalent.

To be more specific, it has been observed that agents tend to take a more aggressive, less compromising stance when under surveillance by their principals. This effect of monitoring may be advantageous for each principal unilaterally, but inefficient overall compared with closed-doors negotiations. What is unclear, however, is why monitoring leads to more aggressive bargaining behaviour in the first place and it is this phenomenon that we seek to explain.

We develop a model in which a seller and a buyer bargain over a good through agents. Our approach to delegated bargaining is very different from that of existing analyses, and is based on three assumptions:

Agents differ in their bargaining skills. Some negotiators are better at ascertaining the opposite side's preferences than others. On average, they negotiate better deals for their principals and are in this sense better bargainers. Accordingly, in our model, we distinguish skilled agents who know the rival principal's reservation price and unskilled agents who do not.

The agents' objective is to promote their reputation as skilled negotiators through their bargaining behaviour. This assumption is based on the observation that agents rarely receive explicit pay for performance. Rather, they stand to get re-appointed if they are perceived to be skilled bargainers.

Principals face uncertainty about their agent's bargaining skill, but also about their opponent's bargaining position (which is why they delegate bargaining in the first place). Hence, while observed bargaining behaviour provides information about an agent, the information is inconclusive. For example, if a union leader negotiates a meager wage increase, workers cannot tell whether the firm's cost situation did not allow a better outcome, or whether their leader was simply a poor bargainer.

With *open-door bargaining*, the principals observe the agents' offers and counter-offers during negotiation, and in particular the final outcome, i.e. trade at some price or disagreement. With *closed-door bargaining*, the principals observe only the final outcome. Our model captures this difference in the simplest possible way: the agents play a one-shot bargaining game in which one agent is randomly chosen to make a price offer and the other accepts or rejects. The principals observe which agent makes an offer if bargaining is conducted with open doors, but only the final outcome with closed doors.

In the equilibrium of this game, a seller infers that their agent is more likely to be skilled if the agent sells at a high price, less likely if they sell at a low price, and least likely if bargaining results in disagreement. Given these beliefs, a skilled agent demands a high price if they know that the buyer's reservation price is high and a low price if they know the reservation price is low. An unskilled agent faces a trade-off: a low price is sure to be accepted by the buyer but is also detrimental to the agent's reputation. A high price, on the other hand, leads to a good reputation if the buyer accepts, but the worst reputation if they reject and bargaining ends with disagreement. In equilibrium, an unskilled agent may randomize between a low and a high price. This sometimes leads to disagreement and hence an inefficient bargaining outcome. Given these strategies, a skilled agent is more likely to sell at a high price and to avoid disagreement, which confirms the principal's beliefs.

Our main result is that with closed-door bargaining, unskilled agents bargain less aggressively than with open doors. Two effects lead to this result. First, since with open doors principals can observe which agent made an offer, while with closed doors they observe only the outcome, closed-door bargaining is less informative about the agents' skills than open-door bargaining. Consequently, with closed doors, selling at a high price results in a less good reputation for the seller's agent and selling at a low price in a less bad reputation, than with open doors.

Second, with both open and closed doors, a principal will blame their own agent if bargaining results in disagreement. The reason is that a skilled agent always trades whenever trade is possible. That is, they bargain aggressively only if they know that the opponent has a weak bargaining position and they can therefore achieve a successful outcome. Otherwise, the agent concedes immediately, knowing that no better outcome can be attained. Disagreement therefore suggests that the agent does not know the opponent's bargaining position, i.e. is unskilled, even when the principal cannot observe the agents' moves during negotiation.

Both effects combined shift the incentives of an unskilled agent in an unambiguous direction: for the seller's agent, offering a low price is less detrimental for their reputation with closed doors than with open ones. Demanding a high price is less attractive, since a successful sale leads to a less good reputation than with open doors, while rejection by the buyer leads to an equally bad reputation. The net effect is that the agent bids a high price less often. Since this also means that disagreement occurs less often, closed-door bargaining is more efficient.

For future bargaining situations, the principals are also interested in learning about their agents' abilities. For any given bargaining strategies, inferences about agents are more precise with open- than with closed-door bargaining. This is offset, however, by the fact that unskilled agents also bargain more aggressively with open doors. We show that the net effect is ambiguous: accounting for the difference in the agents' strategies, open-door bargaining can be more informative overall than closed-door bargaining, or it can be less informative.

Finally, we show that it is not necessarily advantageous for one side to monitor the agents unilaterally when both sides are represented by agents. If the buyer monitors the agents while the seller does not, the buyer's agent bargains more aggressively than with closed doors. As a result, however, the seller's agent is more likely to be credited if trade at a high price occurs, which induces the seller's agent to bargain more aggressively too. Since this reduces the buyer's profit, the buyer may or may not want to monitor the agents. If unilateral monitoring is profitable, however, the seller has an incentive to monitor the agents as well. The principals then face a Prisoners' dilemma where they jointly prefer closed-door bargaining but are led to open-door bargaining by independently choosing to monitor the agents.

1 Introduction

"I. Open covenants of peace, openly arrived at, after which ... diplomacy shall proceed always frankly and in the public view."
Woodrow Wilson: Speech on the Fourteen Points, January 8, 1918

The "process of compromise and negotiation is vital to the adoption of Community legislation, and would be jeopardized if delegations were constantly mindful of the fact that the positions they were taking ... could at any time be made public."
Court of First Instance of the European Communities (1995), rejecting demands to disclose minutes of meetings of the European Council

In most bargaining situations in business and politics, negotiators bargain not for themselves, but on behalf of constituents, or principals. Managers negotiate on behalf of shareholders, union leaders on behalf of workers, politicians on behalf of voters, diplomats on behalf of countries. Two views exist on how closely agents in these situations ought to be monitored by those who appoint them. Some maintain that agents' actions should always be as transparent as possible to the principals represented, and that negotiations should thus be conducted with "open doors". Indeed, some regard openness as a fundamental principle of democracy.[1] Others argue that too much monitoring in some way distracts agents and leads to *less* efficient decision-making processes. In their view, "closed-door" negotiations allow principals to hold their agents accountable for the *results* they negotiate without interfering with the *process* of negotiation.[2]

These two views are reflected in a large variety of institutional arrangements. Labor negotiations often start with a public exchange of demands between managers and union leaders, but are concluded in closed-door meetings. Lawyers prefer to meet with their adversaries in absence of their own clients, while courts sometimes mandate the presence

[1] For example, the decision of the United States Senate to convene behind closed doors on several occasions during the 1999 impeachment trial of President Clinton was criticized by many commentators as well as some senators.

[2] Cf. Fisher and Ury (1981, p.36): *"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"*.

of the principals in settlement negotiations. In parliaments, debates may be televised (as in the U.S.) or held behind closed doors (as in the U.K. until 1989); and voting records may be public (as in Sweden) or secret (as in Germany). Such differences between countries have fueled the recent debate over the appropriate degree of transparency in EU decision-making (see von Sydow 1995). While delegates rountinely emphasize the inefficiencies caused by openness, constituents often suspect that such arguments may be self-serving.

From an agency-theoretic perspective, the benefits of monitoring agents are obvious: monitoring facilitates the rewarding of agents for their performance and thus induces agents to work harder. It can also help prevent collusion between agents. Less obvious is why monitoring can have detrimental effects. The purpose of this paper is to provide an answer to this question, and thus an explanation for why even in democratic institutions closed-door negotiations are prevalent.

To be more specific, it has been observed that agents tend to take a more aggressive, less compromising stance when under surveillance by their principals.[3] As Schelling (1960) conjectured, this effect of monitoring may be advantageous for each principal unilaterally, but inefficient overall compared with closed-doors negotiations. What is unclear, however, is why monitoring leads to more aggressive bargaining behavior in the first place, and it is this phenomenon that we seek to explain.[4]

We develop a model in which a seller and a buyer bargain over a good through agents. Our approach to delegated bargaining is very different from that of existing analyses (which we discuss in Section 8), and is based on three assumptions:

(1) Agents differ in their bargaining skills. Some negotiators are better at ascertaining the opposite side's preferences than others. On average, they negotiate better deals for their principals and are in this sense better bargainers. Accordingly, in our model, we distinguish skilled agents who know the rival principal's reservation price, and unskilled agents who do not.[5]

[3] This effect, pointed out by e.g. Schelling (1960), has been confirmed in many experiments and field studies, cf. Rubin and Brown (1975, p. 43-54) and Pruitt (1981, p. 41-45).

[4] On the other hand, to the extent that delegation allows a principal to *commit* to a certain position (as opposed to making concessions over time), Myerson and Satterthwaite's (1983) efficiency result suggests that delegated bargaining should be more, not less efficient. We discuss this point in detail in Section 8.

[5] In contrast, in previous analyses, agents are assumed to differ in their time discount factor or their

(2) The agents' objective is to promote their reputation as skilled negotiators through their bargaining behavior. This assumption is based on the observation that agents rarely receive explicit pay for performance. Rather, they stand to get re-appointed if they are perceived to be skilled bargainers.[6]

(3) Principals face uncertainty about their agent's bargaining skill, but also about their opponent's bargaining position (which is why they delegate bargaining in the first place).[7] Hence, while observed bargaining behavior provides information about an agent, the information is inconclusive. For example, if a union leader negotiates a meager wage increase, workers cannot tell whether the firm's cost situation did not allow a better outcome, or whether their leader was simply a poor bargainer.

The setup of our model is deliberately one-sided in that it abstracts from any effort agents may have to exert, e.g. to prepare for negotiations. Ignoring obvious reasons in favor of monitoring allows us to focus on the more interesting question of what its drawbacks are.

With *open-door bargaining*, the principals observe the agents' offers and counter-offers during negotiation, and in particular the final outcome, i.e. trade at some price or disagreement. With *closed-door bargaining*, the principals observe only the final outcome. Our model captures this difference in the simplest possible way: the agents play a one-shot bargaining game in which one agent is randomly chosen to make a price offer, and the other accepts or rejects. The principals observe which agent makes an offer if bargaining is conducted with open doors, but only the final outcome with closed doors.[8]

In both scenarios there exists a unique Perfect Bayesian equilibrium that is both

degree of risk aversion, if at all. Cf. Section 8 for a more detailed discussion.

[6] This nature of the principal-agent relationship was already pointed out by Francis Bacon, who, in his essay *Of Negociating*, recommends that principals use agents "such as have been lucky, and prevailed before in things wherein you have employed them; for that breeds confidence, and they will strive to maintain their prescription" (quoted in Nierenberg 1968, p. 121).

[7] Delegation is taken as given in this paper, but is optimal provided that the probability of hiring a skilled agent is sufficiently high. Additionally, in many contexts such as those mentioned above, delegation is not a choice because it is impractical for constituents to meet directly.

[8] We assume throughout that the principals learn the outcome of negotiations ("open covenants"). In practice, this is not always the case, and debates over the pros and cons of open-vs. closed-door negotiations do not always distinguish between secrecy of the outcome and secrecy of the process.

efficient and satisfies a consistency condition (Sections 3 and 4). A seller infers that her agent is more likely to be skilled (compared with her initial beliefs) if he sells at a high price, less likely if he sells at a low price, and least likely if bargaining results in disagreement. Given these beliefs, a skilled agent demands a high price if he knows that the buyer's reservation price is high, and a low price if he knows the reservation price is low. An unskilled agent faces a tradeoff: a low price is sure to be accepted by the buyer but is also detrimental to the agent's reputation. A high price, on the other hand, leads to a good reputation if the buyer accepts, but the worst reputation if she rejects and bargaining ends with disagreement. In equilibrium, an unskilled agent may randomize between a low and a high price. This sometimes leads to disagreement and hence an inefficient bargaining outcome. Given these strategies, a skilled agent is more likely to sell at a high price and to avoid disagreement, which confirms the principal's beliefs.

Our main result (in Section 5) is that with closed-door bargaining, unskilled agents bargain less aggressively than with open doors. Two effects lead to this result. First, since with open doors principals can observe which agent made an offer, while with closed doors they observe only the outcome, closed-door bargaining is less informative about the agents' skills than open-door bargaining. Consequently, with closed doors, selling at a high price results in a less good reputation for the seller's agent, and selling at a low price in a less bad reputation, than with open doors.

Second, with both open and closed doors, a principal will blame her own agent if bargaining results in disagreement. The reason is that a skilled agent always trades whenever trade is possible. That is, he bargains aggressively only if he knows that the opponent has a weak bargaining position and he can therefore achieve a successful outcome. Otherwise, he concedes immediately, knowing that no better outcome can be attained. Disagreement therefore suggests that the agent does not know the opponent's bargaining position, i.e. is unskilled, even when the principal cannot observe the agents' moves during negotiation.[9]

Both effects combined shift the incentives of an unskilled agent in an unambiguous direction: for the seller's agent, offering a low price is less detrimental for his reputation

[9] For the same reason, in more general contexts, the inference is similar if the agent negotiates a poor deal but only after substantial delay. The longer bargaining goes on (which a principal can always observe), the more strongly an agent signals his confidence that the opponent's bargaining position is weak. Accordingly, an eventually poor outcome reflects badly on the agent.

with closed than with open doors. Demanding a high price is less attractive, since a successful sale leads to a less good reputation than with open doors, while rejection by the buyer leads to an equally bad reputation. The net effect is that the agent bids a high price less often. Since this also means that disagreement occurs less often, closed-door bargaining is more efficient.

An implicit assumption of our career-concerns approach is that for future bargaining situations, the principals are also interested in learning their agents' abilities. For any given bargaining strategies, inferences about agents are more precise with open- than with closed-door bargaining. However, unskilled agents also bargain more aggressively with open doors, which distorts the principals' inferences. The net effect is ambiguous, as we show in Section 6. Hence, if open-door bargaining is more informative than closed-door bargaining even accounting for the change in the agents' strategies, the principals face a tradeoff between the efficiency of the current bargaining process and the quality of information and hence the efficiency of future bargaining rounds. If open-door bargaining is less informative overall, the principals unambiguously prefer closed doors.

Finally, we show in Section 7 that, contrary to Schelling's hypothesis, unilateral monitoring is not necessarily advantageous when both sides are represented by agents. If the buyer monitors the agents while the seller does not, the buyer's agent bargains more aggressively than with closed doors. As a result, however, the seller's agent is more likely to be credited if trade at a high price occurs, which induces the seller's agent to bargain more aggressively too. Since this reduces the buyer's profit, the buyer may or may not want to monitor the agents. If unilateral monitoring is profitable, however, the seller has an incentive to monitor the agents as well. The principals then will face a Prisoners' dilemma where they jointly prefer closed-door bargaining but are led to open-door bargaining by independently choosing to monitor the agents.

2 The Model

A seller S and a buyer B employ agents A_s and A_b, respectively, to bargain over an indivisible good. 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 $\underline{s} < \underline{b} \leq \bar{s} < \bar{b}$ (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 $1-p$. Since $\underline{b} < \bar{s}$, trade is not profitable with positive probability.

Figure 1: *Valuations and prices*

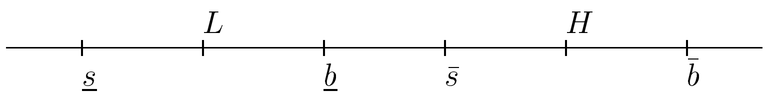

Bargaining is conducted through a simple one-shot game in which with equal probability either A_s or A_b is chosen to make an offer. Only two different prices can be submitted, either high (H) or low (L), where $\underline{s} < L \leq \underline{b} < \bar{s} \leq H < \bar{b}$ (see Figure 1). The other agent accepts, in which case trade takes place at that price, or rejects, in which case no trade takes place and both buyer and seller earn zero payoff. An agent may not submit a bid that can lead to a loss for his principal, but is otherwise not restricted in the choice of his bid. Thus, the agent of a weak seller may offer or accept L or H, while the agent of a tough seller can only offer or accept H, and vice versa for the buyer's agent. We will say that the seller agent is tough if he bids H, and the buyer agent is tough if he bids L. Otherwise, we call an agent weak.

Each agent is hired from a pool of candidates who are "good" with probability q and "bad" with probability $1-q$. Good agents can perfectly observe the other side's reservation price, while bad agents cannot. Neither principal knows her agent's type,[10] and no agent knows the type of the other agent.[11]

The principals have two objectives. One is to maximize their expected payoff from

[10] We have assumed that weak principals give their agents discretion instead of e.g. requiring their agent to always bid one particular price. This assumption can be endogenized: if a principal does not know her agent's type, the benefit of delegation is that a good agent can on average achieve better bargaining results than the principal can if the agent is allowed to condition his bargaining strategy on his knowledge of the rival's reservation price. Accordingly, delegation is optimal if the share of good agents q is sufficiently high. See Appendix B for details.

[11] The idea here is that a good bargainer understands the "fundamentals" of the bargaining situation, i.e. the rival principal's preferences, but does not necessarily know the personal characteristics of his counterpart. The assumption has no consequence in our one-shot model. In a more general context, assuming that e.g. good agents know the other agent's type would affect the analysis. In particular, it would lead to multiple equilibria given that sometime two good agents meet.

bargaining. Another is to obtain information about their agent's ability: for future (unmodeled) bargaining situations, the principals would like to re-appoint agents they believe to be good, or compensate agents according to their perceived ability. We call the probability with which a principal believes that her agent is good the agent's "reputation". Each principal combines the agent's initial reputation q with information gleaned from observed bargaining behavior to form a posterior reputation $\widehat{q}$.[12]

Our assumptions about the agents' objectives follow a career-concerns approach (cf. Holmström 1999). The agents cannot be compensated conditional on the outcome of bargaining, nor do they have any intrinsic interest in the outcome of bargaining. Instead, each receives a wage that depends on his reputation as a bargainer. That is, an agent initially receives a wage proportional to his prior reputation q, whereas in the (unmodeled) future the agent receives a wage proportional to his posterior reputation $\widehat{q}$. For simplicity, we therefore assume that each agent maximizes his reputation as a good bargainer, i.e. maximizes $\widehat{q}$. 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. The main reason for this is that real-world bargaining usually covers multiple issues which often defy measurement.[13]

A principal will be willing to compensate her agent according to perceived ability only if a good agent actually attain a higher payoff for her than a bad one. We therefore

[12] We assume that a principal cannot observe how the other principal reacts toward her agent after bargaining. For example, if union leaders negotiate a mediocre wage increase and shareholders later praise managers for their bargaining skills, we rule out that union members can observe this and infer that the management's position was in fact weak, and that the union leaders therefore must be incompetent.

[13] The assumption that the agents' future payoff is linear in their reputation has two possible interpretations: one is that changes in an agent's reputation with his own principal are reflected, to a lesser extent, in changes in the agent's market reputation and hence wage. Consequently, the principal would then have to pay a wage that changes continuously with his reputation. Alternatively, an agent might receive a fixed wage while his chance of re-appointment depends in a continuous way on his perceived ability. This assumption seems reasonable whenever an agent's bargaining skill is only one of several measures of performance, as is the case in some of our leading examples. In contrast, if the principal has a large pool of potential agents with prior ability q to choose from and reappoints her agent if and only if $\widehat{q} > q$, the agent's payoff function would be discontinuous in his reputation.

restrict the set of equilibria of interest to those satisfying this requirement:

Consistency *(of principals' and agents' objectives): In equilibrium, each principal's expected payoff is higher if her agent is good than if he is bad.* [14]

While S and B are not actively involved in bargaining, they choose an information structure which, via the agents' bargaining behavior, affects their bargaining payoffs and the information they obtain about agents. We compare two scenarios: with *open-door bargaining*, S and B observe the entire bargaining game; i.e. they observe which agent makes an offer, as well as the offer itself and the other agent's response. With *closed-door bargaining*, S and B get to know only the final outcome of bargaining, i.e. the transaction price or disagreement, and use this information to update their beliefs about their agents. In Section 7, we also look at an intermediate case in which B can observe which agent makes an offer but S cannot. The timing of the game is as follows:

1. S and B jointly or independently choose an information structure.
2. S and B hire agents A_s and A_b, respectively. The agents know their own type but the principals do not.
3. The principals' valuations are realized.
4. Each agent learns his own principal's valuation. Good agents also learn the valuation of the other principal.
5. Bargaining: with probability 1/2, either A_s or A_b is chosen to offer a price (L or H). The other agent accepts or rejects.

3 Agents' strategies and payoffs

Strategies: We can simplify the description of the agents' strategies by ignoring the actions that are non-strategic or dominant choices. First, by assumption, agents of weak principals can play either weak or tough, whereas agents of tough principals must play tough. Hence,

[14] This condition can also be seen as a requirement for a principal to want to delegate bargaining to an agent in the first place; for if the condition is violates, a principal can always achieve at least as high a payoff by bargaining for herself than by using an agent.

only agents of weak principals behave strategically. Second, we discard as irrelevant any equilibria that are Pareto-dominated by some other equilibrium for all four players. Applying this criterion, we can restrict our attention to equilibria in which an agent who receives an offer always accepts if the price is acceptable to his principal.[15]

It follows that non-trivial choices are made only by good or bad agents of *weak* principals when *making* price offers.[16] Allowing mixed strategies, an agent's strategy is characterized by the mixing probabilities used by the different types of agents. Let a_i^k be the probability that an agent of type k on side i is tough, where $i = s, b$ denotes the seller or buyer side. Technically, there are three types of agents: a good agent who knows that the other side is tough ($k = t$), a good agent who knows that the other side is weak ($k = w$), and a bad agent ($k = b$). Each agent's strategy can then be described by a triple $\mathbf{a_i} = (a_i^t, a_i^w, a_i^b)$, $i = s, b$, and let $\mathbf{a} = (\mathbf{a_s}, \mathbf{a_b})$.

Events: For open-door bargaining, denote the price quoted by the agent chosen to make an offer by H or L, and the other agent's response by Y or N. Then the set of observable events (with superscript o for open doors) can be described by $E^o = (\{H, L\} \times \{Y, N\}) \cup (\{Y, N\} \times \{H, L\})$, where for example YL denotes the event where A_b offers L and A_s agrees. With closed-door bargaining, the principals observe the final outcome of bargaining but not which agent made an offer. Formally, let $H = \{HY, YH\}$ denote the event that trade at H occurs, $L = \{LY, YL\}$ the event that trade at L occurs, and $\emptyset = \{HN, NH, LN, NL\}$ the event that no trade occurs.[17] The set of observable events with closed-door bargaining then is $E^c = \{H, L, \emptyset\}$.

Payoffs: Let $\widehat{q}_i(e|\mathbf{a})$ be principal i's ($i = s, b$) posterior probability that A_i is good, which depends on the observed event $e \in E^o$ or E^c and both agents' strategies. By assumption, each agent A_i maximizes the expected value of $\widehat{q}_i(e|\mathbf{a})$. From the considerations above, we can restrict our attention to the choices an agent of a weak principal faces when

[15] To see that responding efficiently is not counteracted by the agents' signaling incentives, notice that agents who receive an offer no longer face any uncertainty: if they accept, the good is traded at the price offered, otherwise there is no trade. Since good and bad agents have the same preferences, they must in equilibrium respond identically.

[16] See Section 4 below for a generalization of this point to a model with continuous types.

[17] While we use H and L both for price bids and observable events, there should not be any reason for confusion.

making an offer, and because the model is symmetric, it suffices to look at the seller's side. With open doors, a bid of H can result in the event HY or HN, while a bid of L is always accepted by the buyer and hence leads to the event LY. With closed doors, the corresponding events are H, $\emptyset$ and L.

Denote by $\pi_i(k, p, \mathbf{a})$ $(i = s, b)$ the expected value of the posterior reputation $\widehat{q}(.|\mathbf{a})$ if a type-k agent bids p. If A_s knows that B is weak, then he knows that a price bid of H will be accepted by A_b, whereas a bad A_s can expect A_b to accept H only with probability p. The expected payoffs for each agent type k and each price are given in Table 1 for open-door bargaining. The payoffs for A_b are obtained analogously. The payoffs for closed-door bargaining are as in Table 1, with events H, $\emptyset$ and L in place of HY, HN and LY.

Table 1: Expected payoff $\pi_s^o(k, p, \mathbf{a})$ for A_s

Type	A_s's price bid	
	H	L
w	$\widehat{q}_s(HY\|\mathbf{a})$	$\widehat{q}_s(LY\|\mathbf{a})$
d	$p\widehat{q}_s(HY\|\mathbf{a}) + (1-p)\widehat{q}_s(HN\|\mathbf{a})$	$\widehat{q}_s(LY\|\mathbf{a})$
t	$\widehat{q}_s(HN\|\mathbf{a})$	$\widehat{q}_s(LY\|\mathbf{a})$

4 Equilibrium

Our model has features of both a cheap-talk and a signaling game: bids are costless and affect an agent's payoff only indirectly through his principal's inferences. As a result, the game has many equilibria. On the other hand, good and bad agents have the same preferences over the principal's beliefs: both would like to perceived as good. In a normal cheap-talk game, these preferences would render communication meaningless. Here, however, the agents nevertheless separate statistically because of the good types' better knowledge about the opponent.

A *Perfect Bayesian Equilibrium* in the bargaining game with reputation concerns consists of a pair of strategies $\mathbf{a}$ and beliefs $\widehat{q}_i^o(e|\mathbf{a})$ such that

(1) each agent type's action at stage 5 of the game is optimal given the strategies of all

other types and of the rival agent, and given the principals' beliefs, i.e.

$$\begin{aligned} a_s^k &= \arg\max_{a\in[0,1]} a\pi_s(k,H,\mathbf{a_s^k}(a)) + (1-a)\pi_s(k,L,\mathbf{a_s^k}(a)) \quad \text{and} \\ a_b^k &= \arg\max_{a\in[0,1]} a\pi_b(k,L,\mathbf{a_b^k}(a)) + (1-a)\pi_b(k,H,\mathbf{a_b^k}(a)) \quad \text{for} k=t,w,d \end{aligned} \tag{1}$$

where $\mathbf{a_i^k}(a)$ denotes the strategy combination $\mathbf{a}$ with a_i^k replaced with a; and

(2) the principals' beliefs $\widehat{q}_i^o(e|\mathbf{a})$ satisfy Bayes' rule for any event e that occurs with positive probability, i.e.

$$\widehat{q}_i(e) = \frac{\Pr_i(e|\text{good})q}{\Pr_i(e|\text{good})q + \Pr_i(e|\text{bad})(1-q)}, \tag{2}$$

where $\Pr_i(e|k)$ is the probability that an agent of type $k \in \{good, bad\}$ reaches event e, evaluated from principal i's point of view.[18]

We then have

Proposition 1 *With both open- and closed-door bargaining, there exists a unique Perfect Bayesian Equilibrium that is efficint and satisfies the Consistency condition. In this equilibrium, (a) a good agent who knows the rival is weak plays tough ($a_i^w = 1$), (b) a good agent who knows the rival is tough plays weak ($a_i^t = 0$), and (c) a bad agent randomizes between the two prices or plays weak. Specifically,*

- *with open doors, a bad agent plays tough with probability*

$$a^o := \frac{p-q+(1-p)pq}{1-q} < p \tag{3}$$

$$\text{if} \quad \frac{p}{(1-p)^2} > \frac{q}{1-q} \tag{4}$$

(a sufficient condition is $p > q$), and otherwise he plays weak ($a_i^b = 0$).

- *With closed doors, a bad agent plays tough with probability*

$$a^c = \frac{\sqrt{p}[5-q-p(3-q)] + \sqrt{4(2-p)-p(1-p)(1-q)[3+q+p(1-q)}}{2\sqrt{p}(1-p)(1-q)} \tag{5}$$

$$\text{if} \quad p(3-p)(2-q) < 2, \tag{6}$$

and otherwise he plays weak ($a_i^b = 0$).

[18] Again, there should be no confusion when we index the agents' types by $k \in \{good, bad\}$ when looking at a principal's inferences but by $k \in \{t, w, b\}$ when discussing the agents' strategies.

S's corresponding posterior beliefs satisfy $\widehat{q}^o_s(HY) > q > \widehat{q}^o_s(LY) > \widehat{q}^o_s(HN)$ with open-door bargaining; and analogous conditions hold for B's beliefs and for closed doors.

All proofs are in the Appendix. Consider the case of open-door bargaining; the intuition for closed doors is the same. If $\widehat{q}$ is increasing in the price at which A_s trades, i.e. $\widehat{q}(HY) > \widehat{q}(LY) > \widehat{q}(HN)$, then a good A_s who knows that B is weak will play H since $\widehat{q}(HY) > \widehat{q}(LY)$, whereas a good A_s who knows that B is tough will play L since $\widehat{q}(LY) > \widehat{q}(HN)$. For a bad agent, in contrast, playing H is a gamble that can lead to either the best or the worst outcome (trade at H or disagreement), while playing L is a safe option. If p, the probability that B is weak, is sufficiently high, then in equilibrium a bad A_s mixes between playing H and L, whereas if p is smaller, he always plays L.

Mixing can occur because the principal's posterior beliefs vary with the agent's strategy: the more often a bad A_s plays tough, the less S believes to have a good agent upon observing HY, and the more she believes to have a good agent upon observing LY. Hence $\widehat{q}(HY)$ is decreasing in a^b_s and $\widehat{q}(LY)$ is increasing in a^b_s. Thus, if for some a^b_s a bad agent prefers H to L because $p\widehat{q}(HY|\mathbf{a}) > \widehat{q}(LY|\mathbf{a})$, an increase in a^b also reduces the difference between the expected payoffs of playing H or L. An equilibrium is then reached at an a^b for which the expected payoffs of H and L are equal.

In equilibrium, observing that A_s sells at L must be bad news about him, i.e. $\widehat{q}(LY) < q$: if the prior probability that A_s is good is q, the posterior associated with an event that is a safe option available to all agents cannot exceed q, and must be strictly smaller if good agents sometimes choose a different option. This in turn means that a bad agent must play L at least as often as a good agent, i.e. $1 - a^b > 1 - p$ or $a^b < p$.

Since the agents' price bids are cheap talk and are chosen only because of the meaning the principals attach to them, more equilibria exist. With open-door bargaining, another separating equilibrium exists in which A_s's perceived skill is *decreasing* in the negotiated price, i.e $\widehat{q}(HY) < \widehat{q}(LY) < \widehat{q}(HN)$. Facing these beliefs, a good agent bids H if he knows B is tough and L if B is weak, while a bad agent bidding H may sometimes end up at the "worst" outcome, trade at H. Since again good agents are on average more successful in attaining whatever is interpreted as indicative of bargaining skill (here, disagreement), this confirms the principal's beliefs.[19] However, this equilibrium violates the Consistency

[19] With closed doors, this equilibrium does not exist: if disagreement is interpreted as good news, any

criterion introduced above: since good agents systematically negotiate *worse* outcomes for their principals than do bad agents, the principals would not want to pay a higher wage for good agents (on the contrary), in which case the agents would have no reason to signal that they are good. We therefore do not consider this equilibrium further. Additionally, because bids are cheap talk, any combination of pooling strategies, together with beliefs that ignore which event is observed, i.e. $\widehat{q}_i(e) = q$ for all e, constitutes a ("babbling") equilibrium. These, too, violate Consistency.[20]

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_i^b, or simply a_i, of a bad agent. Abusing the notation, let $\mathbf{a} = (a_s, a_b)$.

Notice that the equilibrium described here is a result of the agents' incentives to signal their knowledge about *something*, and does not depend on a bargaining context as such. For example, suppose a firm hires an agent to forecast demand for its product. Good agents know the level of demand, bad agents do not. Suppose that correctly predicting high demand leads to a higher profit for the firm than correctly predicting low demand, but that profits are lowest if demand is low but the agent predicts high demand. If the agent's objective is to maximize his reputation, the situation is equivalent to the one considered above, and the resulting equilibrium is as described in Proposition 1.

Digression: an extension to continuous types and prices

To illustrate how some features of our simple model carry over to a continuous framework, consider a model with one-shot sequential bargaining as in our main model but with continuous types and prices as in Chatterjee and Samuelson (1983) and Myerson and Satterthwaite (1983), see Appendix D for details. Suppose both the seller's and the buyer's reservation prices are uniformly distributed over the unit interval, and that agents can bid any price on the unit interval. If trade is profitable, good agents simply bid the rival's reservation price, knowing that it will be accepted. Otherwise, they announce that

agent receiving an offer would always reject, knowing that his principal cannot tell whether he made or received an offer.

[20] Pooling equilibria would also be eliminated by assuming that agents receive some ε-share of the principals' payoffs. An arbitrarily small share, however, would not suffice to eliminate the "reverse" separating equilibrium described above.

trade is not possible. Since a pure strategy would give away a bad agent, bad agents must in equilibrium randomize among the prices acceptable to their principal (including the announcement that trade is not possible). An equilibrium is then characterized by a function that maps each principal's type into a distribution of prices bid by a bad agent.

A separating equilibrium similar to that of Proposition 1 can be explicitly determined for open-door bargaining but not for closed-door bargaining. It has three main features. First, disagreement always leads to the worst reputation even if trade is not always possible. The reason is that a good agent always reaches an agreement whenever it is possible, while a bad agent does not.

Second, the higher the price a bad seller agent bids, the larger is the risk of rejection. For this reason, bad agents bid high prices less frequently than lower prices, while good agents bid high and low prices equally often (because of the uniform distribution of the buyer's valuation). Consequently, if trade takes place, the agent's resulting reputation is increasing in the price at which he trades.

Third, the seller agent's reputation if no trade occurs is continuously increasing in the seller's reservation price: if the reservation price is zero, then failure to trade is a certain sign that the agent is bad. However, if it is high, then with high probability even good agents are unable to trade, and if trade occurs, it must be at a high price regardless of the agent's ability. Therefore, a seller with a high reservation price on average learns little about her agent whether or not trade occurs. This is the continuous analog of a stark feature of our two-type model, namely that bids and outcomes are informative of an agent's type if his principal is weak but not if she is tough (since in the latter case any agent is required to be tough). Quite generally, agents of tough principals have little freedom in their bargaining behavior, and accordingly their principals learn little about them. Meaningful decisions and inferences arise only for agents of weak principals.

5 Closed-door bargaining is more efficient

Our main result, a prediction about the difference in the bargaining behavior of bad agents, is based on a comparison of the posterior beliefs associated with observable events in the two scenarios. We first derive this result in the context of our model, and then discuss how it extends to a two-stage version of our bargaining game.

5.1 The main result

Since A_s's price bids reveal information about his type but his responses do not, trade at any price is less informative with closed than with open doors because S cannot see which agent made an offer. That is, with closed doors trade at a high price is good news about A_s, but not as good as with open doors because S cannot see whether her own agent or A_b offered a high price. Similarly, A_s's reputation if he trades at L is below q, but still better than with open doors.

Upon observing no trade ($\emptyset$), however, a weak seller's inferences are the *same* for closed doors as for open doors: her agent must be bad, i.e. $\widehat{q}(\emptyset) = 0$. The reason is simply that a good agent always trades. This inference is based in part on the seller's knowledge that since she is weak, any A_s would accept any offer from A_b. Disagreement therefore must mean that A_s made an offer. Formally, we have

Lemma 1 *If $a_s, a_b < p$, then $\widehat{q}(HY) > \widehat{q}(H)$, $\widehat{q}(LY) < \widehat{q}(L)$, and $\widehat{q}(HN) = \widehat{q}(\emptyset) = 0$.*

Lemma 1 is illustrated in Figure 2. The parameter constraint on a_s and a_b is a requirement that follows from the equilibrium conditions $\widehat{q}(L) < q$ and $\widehat{q}(LY) < q$ in Proposition 1. The formal proof that a^c in (5) is indeed smaller than p is a by-product of Proposition 2 below.

Figure 2: *Posterior beliefs with open-and closed-door bargaining for* given *strategies*

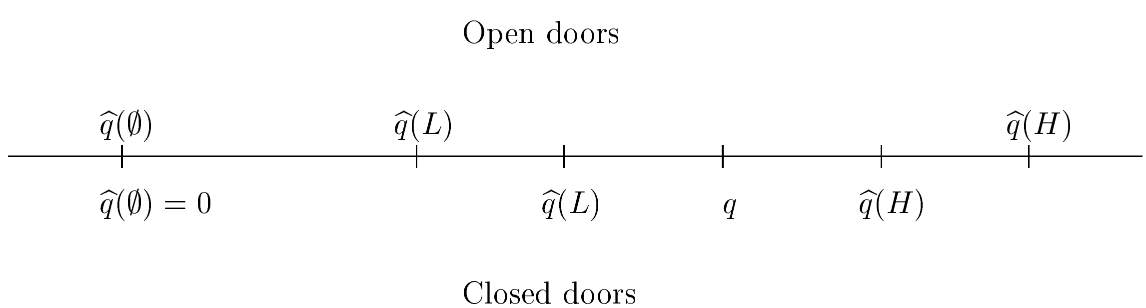

Our main result is the following:

Proposition 2 *With closed doors, a bad agent of a weak seller plays H less often than with open doors, for any given strategy a_b. Therefore, in equilibrium, bad agents of both* *sides bargain less aggressively with closed doors, i.e.* $a^c \leq a^o$. *Hence, the expected surplus from bargaining is larger with closed doors than with open doors.*

To see why, compare the posteriors $\widehat{q}$ for open- and closed-door bargaining (see also Figure 2). According to Lemma 1, we have for any given a_s and a_b:

1. $q < \widehat{q}(H) < \widehat{q}(HY)$ and $\widehat{q}(\emptyset) = \widehat{q}(HN)$, implying that playing H is unambiguously less attractive with closed-door bargaining than with open doors, and

2. $q > \widehat{q}(L) < \widehat{q}(LY)$, implying that playing L is unambiguously more attractive with closed-door bargaining.

Thus, unless bad agents already always play weak with open-door bargaining, switching to closed doors will lead them to reduce the probability of playing tough. This can lead either to an equilibrium in which the agents again mix (because the principals beliefs adjust as the agents change their behavior), or always play weak. Since expected total surplus is decreasing in the bad agents' probability of playing tough, bargaining behind closed doors is more efficient.[21]

While demonstrated for a very simple and specific bargaining model, the two effects that drive this result, and hence the result itself, appear to be general, which we illustrate in Section 5.2 by looking at an extension of our model to two bargaining rounds.[22]

First, closed-door bargaining is generally less informative about the agents' skills than open-door bargaining because the principals cannot observe the sequence of offers and counter-offers that lead to any given outcome. This information is relevant for a principal's inferences about her agent: since any offer made by the opponent reveals information

[21] According to Proposition 2, A_s bargains less aggressively behind closed doors for any given strategy a_b used by a bad A_b. This result is likely to extend to any buyer strategy (a_b^t, a_b^w, a_b^b), in which case the first part of Proposition 2 would not depend on the presence of an agent on the buyer side; it could also be B herself. In our one-shot game, B's optimal strategy in this case would not depend on A_s, for she is only concerned with the payoff from bargaining, which does not depend on A_s if B is make an offer. We have not analyzed, however, what a closed-door equilibrium between A_s and B might look like in a more general bargaining game such as the one analyzed in Section 5.2.

[22] In addition, the same effects and results are obtained in a bargaining game with simultaneous bids, although such a model has other features that make it less attractive than a sequential model. See Appendix E for details.

about the opponent's bargaining position, the own agent's information about the opponent, i.e. his skill, is less important when he responds to an offer than when he makes one. Accordingly, with open doors, offers the own agent makes are more informative of the own agent's skill than responses to the opponent's offers. With closed doors, this distinction cannot be made. Hence, while trade at a high price is still good news and trade at a low price bad news about the agent, the beliefs are more *compressed* toward the prior than the corresponding open-door beliefs.

This compression effect does not apply when bargaining results in disagreement, or more generally, when the agent negotiates a poor deal but only after substantial delay. Here, a principal with a weak bargaining position will tend to blame her own agent with closed- as well as with open-door bargaining, i.e. her beliefs are *anchored*. The reason is that by choosing whether to bargain aggressively or to be appeasing, a bad agent essentially decides between mimicking a good agent who knows the opponent has a weak bargaining position, and a good agent who knows the opponent has a strong bargaining position. The agent's success in mimicking either type depends on whether the resulting outcome is one that could have been reached by a good agent. But a good agent bargains aggressively only when he knows that the rival's bargaining position is weak and that he can therefore (eventually) achieve a successful outcome. Also, a principal with a weak bargaining position knows that failure to trade cannot result from stubbornness of the rival's agent alone, for her own agent would in this case eventually concede. Hence, disagreement or delay without eventual success suggest to a principal that her own agent misjudged the rival's position even when she cannot observe the agents' moves during negotiation.[23]

Both effects combined shift the incentives of a bad agent in an unambiguous direction as one moves from open- to closed-door negotiations: appeasing bargaining behavior is

[23] In our model, the inference that the own agent is bad if disagreement occurs is perfect, as explained above. This inference remains perfect in a multi-stage bargaining game because good agents of a principal with a weak bargaining position always reach some agreement. The inference is less clear-cut if the agent reaches a bad agreement after some delay: It possible that a skilled agent who knows the rival's position is weak nevertheless fails to reach a favorable outcome because of the rival agent's stubbornness. However, the longer negotiations continue, the less likely this outcome becomes, so that, once again, a principal will tend to blame a poor outcome on her own agent.

less detrimental for the reputation of a bad agent because own concessions are less visible. Aggressive bargaining behavior is less attractive: if it succeeds, the resulting reputation is less good than with open doors because the principal cannot distinguish the own agent's toughness from the rival's appeasing behavior. On the other hand, if it fails, the resulting delay or disagreement are hardly less detrimental for the agent's reputation than with open doors. Consequently, the agent prefers to bargain less aggressively.

5.2 Two-stage bargaining

To illustrate some of the effects that arise when bargaining is conducted in multiple rounds, consider a model in which the bargaining game of Section 2 is played twice (see Appendix C) for details. That is, if the first offer by one agent is rejected by the other, then another round is played in which again one agent is randomly chosen to make an offer. The sets of observable events and the agents' strategy spaces are now much larger than in the one-shot game. In particular, accepting or rejecting an offer in the first round now is a strategic decision.

We determine an equilibrium in which the seller agent's reputation is increasing in the price at which he trades, as above. A good agent then always bids a high price, and rejects a low price in the first round, if he knows the buyer has a high reservation price, but concedes immediately if the buyer's reservation price is low. A bad agent randomizes between a high and low price in the first and possibly the second round, and also randomizes between accepting and rejecting a low price offered by the buyer in the first round. We can numerically compute the symmetric equilibrium of this game for both open and closed doors. A comparison leads to the result of Proposition 2: in both rounds, a bad agent bids a high price more often with open than with closed doors.

The same forces as in the one-shot model are at work here: by choosing whether to bid high or low and whether to accept or reject a low price in the first period, a bad agent decides to mimic one or the other type of good agent. In particular, with open doors, if he chooses to bid high he must continue to bid high in the second period even if the buyer's agent rejects, for being tough initially but appeasing later would be certain evidence that he is bad. Accepting a low price in the second round, in contrast, does not convey much information about the agent's type, because in the last round of bargaining both types of agents accept any offer that their principal can accept.

If bargaining is behind closed doors instead, a bad agent's incentives change in the same way as above. For agreements reached in the first round, the principal's beliefs are compressed toward the prior: selling at a high price leads to a less good reputation than with open doors because the high price could have been offered by the buyer. Similarly, selling at a low price leads to a less bad reputation. Both effects encourage more appeasing behavior in the first round.

When bargaining continues to the second round, a bad agent effectively signals to his principal that he believes the buyer's position to be weak, even if the principal cannot observe the course of bargaining. Consequently, selling at a high price leads to a good reputation even if it is the buyer who offers the high price: a good outcome combined with the agent's prediction (in the first round) that the rival is weak are good news about the agent. Disagreement, on the other hand, leads to the same inference as with open doors: since a good agent always reaches an agreement, the agent must be bad. Thus, bidding a high price in the second round leads to the same reputation, compared with open doors, whether the bid is successful or not. But with closed doors, bidding a low price is now an option because the seller can't see which agent makes a bid. This option is additionally attractive because the outcome of the first round signals to the agent that the buyer might be tough after all. In the closed-doors equilibrium, a bad agent is no longer locked into an aggressive bargaining strategy in the second round, and randomizes between a high and a low price.

5.3 Reputation as viewed by outsiders

We have assumed that agents maximize their reputation as perceived by their own principals. Alternatively, agents might be concerned with their reputation as perceived by an external market that observes the process or only the outcome of negotiations, but not the reservation prices of the principals involved. It turns out that in this case our main result no longer holds; the agents' behavior with open and with closed doors is only insignificantly different.

A seller with low reservation price infers that her own agent is bad if no trade takes place. This inference is based on the knowledge that a good agent would have agreed to a low price if he knew that the buyer's reservation price was low. The inference is hence implicitly based on the seller's private information about her reservation price. A seller with high reservation price, in contrast, would not blame her agent for failing to reach an agreement. Since an external observer cannot make this distinction, the anchoring effect vanishes. Inspection shows that with closed doors, all posteriors $\widehat{q}(e)$, including $\widehat{q}(\emptyset)$, are compressed towards to the prior. Hence, a bad agent's response to a change in the information structure no longer has a predictable direction.

We conclude that the assumption that agents want to leave a good impression on their current principals is an important part of our story. This seems to match well with the applications we have in mind: politicians, diplomats, and union leaders engage in posturing when observed because they have an interest in being re-appointed by their own organization (their party, country, union) and not a different one. This extends to the cases in which an agent's market reputation depends on the word-of-mouth of former clients rather than on what the market literally observes. By contrast, if an agent's market reputation rests more on what the market can observe (for instance, attorneys' observed performance in court), posturing because of monitoring is less likely to occur.

6 Information about agents

Monitoring the agents leads to a less efficient bargaining outcome but may at the same time provide the principals with more accurate information about their agents' skills. It is intuitively clear that for *given* bargaining strategies, the information that principals gain about their agents must be more precise with open-door bargaining. However, bad agents bargain more aggressively with open doors, which further distorts the principals' inferences. We now show that the net effect is ambiguous: open-door bargaining can be more informative than closed-door bargaining, or less informative.

To describe the informativeness of the two scenarios formally, denote an agent's true ability by q_0, which takes the value 1 for a good agent (with probability q) and 0 for a bad agent. The principal's belief $\widehat{q}(e)$ can then be regarded as a point estimator of the agent's true ability, where as before we adopt the seller's perspective. Define $u = \widehat{q}(e) - q_0$ as the estimation error. When the agent is bad (i.e. $q_0 = 0$) and the principal's belief is $\widehat{q}(e)$, we have $u = \widehat{q}(e)$. If the agent is good ($q_0 = 1$), then $u = \widehat{q}(e) - 1$. It is straightforward to show that $\widehat{q}$ is an unbiased estimator of q_0, i.e. that u has a mean of zero.

We use the variance of u as a measure of the noisiness of the information structure: the smaller $\mathrm{Var}(u)$, the more precise the principals' information about their agents. Write $\Pr(e,k) = \Pr(e|k)\Pr(k)$ to denote the joint probability that the agent is of type $k \in \{\text{good, bad}\}$ and event e is observed. Then we have

$$\mathrm{Var}(u) = \mathrm{Var}(\widehat{q} - q_0) = \sum_e \left\{ [\widehat{q}(e)]^2 \Pr(e, \text{bad}) + [\widehat{q}(e) - 1]^2 \Pr(e, \text{good}) \right\}.$$

The next result provides a much simpler expression:

Lemma 2

$$Var(u) = \sum_e \Pr(e, bad)\widehat{q}(e) = (1-q)\sum_e \Pr(e|bad)\widehat{q}(e).$$

Using Lemma 2, we can first show that

Proposition 3 *For any given strategies* $\mathbf{a}$*,* $Var(u)$ *is larger for closed-door than for open-door bargaining.*

The result is straightforward because the set of observable events is strictly coarser for closed- than for open-door bargaining. Thus, holding the agents' bargaining strategies fixed, open-door bargaining is more informative. This is no longer generally true if we compare the variances of estimation errors of the open- and closed door *equilibria*. To see this, we calculate $\mathrm{Var}(u)$ for both open and closed doors, using Lemma 2. Half of the time A_b makes an offer, which A_s always accepts. With open doors, S can see this and the resulting posterior reputation for A_s is q. If a bad A_s makes an offer, he bids H with probability a^o, which is accepted with probability p and then leads to the posterior $\widehat{q}(HY)$. Otherwise he bids L and gains the reputation $\widehat{q}(LY)$. Thus we have

$$\begin{aligned} \mathrm{Var}(u^o)/(1-q) &= \sum_{e \in E^o} \Pr(e|\text{bad})\widehat{q}(e) \\ &= \frac{1}{2}q + \frac{1}{2}\left[a^o p \widehat{q}(HY) + (1-a^o)\widehat{q}(LY)\right]. \end{aligned} \tag{7}$$

With closed doors, trade at any given price leads to the same reputation regardless of whether A_s or A_b makes an offer. If S is weak, only a bad A_b will ever bid H, hence the probability of that A_b bids H is $p(1-q)(1-a^c)$, and otherwise A_b bids L. If A_s makes an offer, he arrives at H with probability pa^c and at L with probability $1 - a^c$, whereas if he bids H and A_b rejects, the result is $\widehat{q}(\emptyset) = 0$. Then we have

$$\begin{aligned}
\mathrm{Var}(u^c)/(1-q) &= \sum_{e \in E^c} \Pr(e|\mathrm{bad})\widehat{q}(e) \\
&= \frac{1}{2}\{[1 - p(1-q)(1-a^c)]\widehat{q}(L) + p(1-q)(1-a^c)\widehat{q}(H)\} + \\
&\quad \frac{1}{2}[a^c p\widehat{q}(H) + (1-a^c)\widehat{q}(L)] \\
&= \{1 + [1 - p(1-q)](1-a^c)\}\widehat{q}(L) + p[1 - q(1-a^c)]\widehat{q}(H). \qquad (8)
\end{aligned}$$

By example we can show

Proposition 4 *The variance of the estimation error $u = \widehat{q} - q_0$ can be larger or smaller in the closed-door equilibrium than in the open-door equilibrium.*

Thus, two cases can arise: if open-door bargaining is more informative about the agents' abilities in equilibrium, the principals face a tradeoff between the efficiency of the current bargaining situation and the efficiency of future bargaining situations. However, it may also be that the advantages of transparency in open-door bargaining are entirely undone by the agents' more aggressive bargaining. In this case, the principals unambiguously prefer closed doors.

7 Unilateral Monitoring

We have so far shown that principals jointly prefer to let agents negotiate secretly if efficiency is the primary concern. Suppose now that it is possible for each side to decide unilaterally whether to monitor the agents. Imagine, say, 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?

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 1 and does not depend on A_s's strategy. Since $a^o \geq a^c$, A_b bargains more aggressively than behind closed doors if monitored by B. S in contrast faces the same situation as with closed doors; hence A_s's best response is given by $a_s^*(a^o)$ according to (16). Intuition may suggest that A_s responds to an increase in a_b by reducing a_s. This might be true if S negotiated with A_b directly: since a mediocre agreement is better than none at all, it might be best to be appeasing if the rival is aggressive.[24] This argument does not apply, however, if S is represented by an agent concerned with his reputation and not his principal's payoff:

Proposition 5 *If B monitors the agents and S does not, there exists a unique Perfect Bayesian equilibrium that is efficient and satisfies the Consistency criterion. In this equilibrium, good agents use pure strategies as described in Proposition 1. Bad agents play tough with probabilities $a_b = a^o$ and $a_s \in (a^c, a^o)$, respectively.*

The proof shows that A_s's closed-door reaction curve in (a_s, a_b)-space is upward-sloping if a_b is greater than or equal to a^c. Thus, A_s responds to an increase in a_b by increasing a_s as well. To see why, look at how the beliefs $\widehat{q}(H)$ and $\widehat{q}(L)$ depend on a_b. Both are increasing in a_b for different reasons. An increase in a_b leads to an increase in $\widehat{q}(H)$ because if A_b is more likely to bid L, then trade at H is more likely to be brought about by A_s. A higher a_b also raises the chances that trade at L is brought about by A_b. This increases the "dilution" of $\widehat{q}(L)$ due to offers made by A_b and moves it closer to the prior q. Since $\widehat{q}(L) < q$, this means that $\widehat{q}(L)$ must increase as well. In our model, the first effect always exceeds the second, inducing A_s to increase a_s.

Proposition 5 suggests that the incentives for B to monitor the agents unilaterally are ambiguous. First, it depends on the parameters of the model whether a weak B *wants* a bad A_b to be tough or not: bidding L leads to trade at L with probability p and hence to an expected payoff of $p(\bar{b} - L)$; bidding H leads to a payoff of $\bar{b} - H$. Thus B wants A_b to bid L if

$$p(\bar{b} - L) > \bar{b} - H, \tag{9}$$

i.e. if the probability that S is weak is sufficiently large, otherwise B would prefer her agent to bid H.

[24] In our model, a game with direct bargaining between a weak S and B has two pure-strategy equilibria, where either both bid L or both bid H. These equilibria lead to asymmetric payoffs for the principals. The only symmetric equilibrium involves mixing between L and H, in which case the expected total surplus is lower than in either pure-strategy equilibrium.

Second, B's payoff is always decreasing in a_s, which according to Proposition 2 increases when B monitors. Thus, (9) is a necessary but not sufficient condition for unilateral monitoring to be profitable. Simulations show that the increase in a_s is small relative to the increase in a_b. Consequently, if $p(\bar{b}-L)$ is sufficiently larger than $\bar{b}-H$, monitoring is profitable for B even if A_s responds by bargaining more aggressively. If B does monitor, however, then S has an incentive to reciprocate by monitoring too: if the analog of (9) holds for S, she would like her agent to be tough as well. If B already monitors and hence $a_b = a^o$, there is no further adjustment of a_b if S monitors, hence S would definitely want to monitor. Hence, if B has only a slight preference for a bad A_b to be tough, she would lose by monitoring because of A_s's response. Formally, we have

Proposition 6 *Unilateral monitoring can be profitable for B only if (9) holds. If (9) holds and B monitors unilaterally, and if the principals' reservation prices are symmetric ($\bar{b} = 1 - \underline{s}$ and $\underline{b} = 1 - \bar{s}$), then S has an incentive to monitor as well.*

To conclude, if both sides negotiate through agents, unilateral monitoring may not be profitable. If it is, though, the principals face a Prisoners' Dilemma, as first suggested by Schelling (1960): they may jointly prefer bargaining behind closed doors, but end up with open-door bargaining because of their individual incentives to monitor.

8 Related literature

Delegated bargaining was first discussed by Schelling (1960), who argued that delegation allows a principal to credibly commit to some bargaining position and thereby gain an advantage. He also suggested that the bilateral use of this tactic could lead to an inefficient outcome. These hypotheses have been analyzed in various forms in the literature.

A large literature has studied the possibilities for a principal to commit to some strategy through a contract with an agent. An important general finding of this literature is that commitment through a contract is credible only if secret renegotiation of the contract is costly, or if the agent's objectives are genuinely different from the principal's (see Katz, 1991, and Caillaud and Rey, 1994).[25] The case of costly renegotiation is studied

[25] Conditions under which delegation can work even when contracts are observable are studied in

by Bester and Sakovics (1998), whereas most other analyses rely on differences between a principal's and her agent's objectives.

For example, one strand of the literature (Jones 1989, Burtraw 1993, Segendorff 1998) studies games in which principals hire agents who subsequently agree on the Nash Bargaining solution. In these models each principal individually has an incentive to appoint an agent whose preferences (e.g. degree of risk aversion), via the Nash Bargaining solution, confer an advantage upon the principal. The bilateral use of this opportunity, however, can lead to an outcome that is inefficient compared with bargaining directly between the principals.

While the conditions for the feasibility of commitment through delegation have been studied extensively, less appears to be known about what consequences the bilateral use of commitment tactics has. In a context of strategic bargaining in multiple stages, a seller who is incompletely informed about the buyer's reservation price is always better off if he can commit to a take some offer instead of offering a sequence of prices that decrease over time (see Kennan and Wilson 1993, p. 58). This situation is equivalent to that of a durable-goods monopolist.

It is not true, however, that the bilateral use of commitments leads to greater inefficiency than bargaining without commitments, on the contrary. Myerson and Satterthwaite (1983) showed that the ex-ante efficient mechanism in a double auction with two-sided incomplete information is one in which some types of sellers and buyers always trade, and others never. Cramton (1985) derived this result in a context of bargaining over time: in an ex-ante efficient incentive compatible mechanism, trade takes place either immediately or never, but not at some later point in time. Intuitively, if commitments are not feasible and the parties know that negotiations will continue if offers have been rejected, then initial demands will be high and early agreement unlikely.[26] Bilateral commitments avoid this Coasian dynamics by forcing the parties to make "reasonable" take-it-or-leave-it offers. Thus, while incomplete information gives rise to inefficiency, the ability for both parties to commit to a bargaining position tends to lead to an increase,

Feshtman and Kalai (1997).

[26] This result may also explain the experimental finding of Schotter et al. (2000) that bargaining through agents was more likely to end in disagreement if the agents and principals could consult several times during bargaining than if agents received instructions only once at the beginning.

not a decrease, in expected surplus.[27]

This result poses a puzzle: if the effect of monitoring is to facilitate the use of commitment tactics, Myerson and Satterthwaite's result suggests that the outcome should be more efficient than bargaining without commitments, i.e. with closed doors, contrary to intuition. Our theory resolves this puzzle by linking the distinction of open vs. closed doors not with commitments in a literal sense but with the agents' reputation concerns. In our theory, "posturing" results not because principals instruct their agents to be tough; rather, the agents choose to be tough in order to distort their principal's information about their skill, and it is this form of posturing that leads to inefficiency.

In this respect, Cai's (2000) model is similar to ours. He analyzes a model in which one side is represented by an agent, who incurs private costs of bargaining. Potential agents differ in these costs. The constituents prefer to have a low-cost agent who doesn't easily concede to the other side, but they face uncertainty about both their agent's type and the other side's reservation value. Patient agents then have an incentive to signal their type by delaying negotiations, which is inefficient.

While the agents' signaling incentives arise in our paper for similar reasons as in Cai's, there are three major differences between the papers. The first is in focus: Cai's emphasis is on the loss of efficiency due to delegation because of the agent's signaling, whereas our focus is the comparison of informational environments within a similar (two-sided) signaling context. Second, Cai studies incentive contracts, whereas we use a career-concerns approach. Third and most importantly, when agents differ in their time preference, principals should be able to pay the more impatient types to act as patient bargainers, which Cai rules out. In contrast, in our model agents differ in their ability to ascertain the rival's reservation price, and it is impossible for principals to pay unskilled agents to be skilled.

Perry and Samuelson (1994) analyze open- vs. closed-door bargaining from an angle

[27] Crawford (1982), too, emphasizes that an increase in the difficulty of commitment does not generally reduce the probability of impasse, contrary to what is commonly assumed. The results discussed here also relate to the objection that since the outcomes of finite-stage games such as ours are sometimes ex-post inefficient, there is scope for renegotiation, which is not allowed. The response to this objection, however, is that allowing renegotiation can only make matters worse, i.e. reduce expected total surplus. Thus, models that rule out renegotiation do not add an inefficiency that might not exist in real-world bargaining, but in fact assume that circumstances are more favorable than they might be in reality.

quite different from ours. They examine bargaining between two parties one of which is represented by an agent. Under closed-door bargaining, heterogeneous constituents must approve the final bargaining agreement. With open doors, constituents may also terminate bargaining after initial offers have been made and rejected. Perry and Samuelson show that if the population of constituents changes over time, early termination in the open-door scenario can result if constituents fear that future voters might approve of an agreement that makes them worse off than the status quo. This threat of termination induces the rival to make a larger concession in the first round, implying that open-door bargaining is advantageous for the side with the agent.

Outside the world of bargaining models, there is of course a larger literature that analyzes the drawbacks of monitoring an agent. In Aghion and Tirole (1998), monitoring and intervention by the principal discourages the agent from investing effort in the evaluation of suitable projects. In Crémer (1994), monitoring that leads to positive information about the agent's ability makes it difficult for a principal to threaten to fire her agent if performance is bad. In Friebel and Raith (2000), intra-firm communication can lead a principal to replace her manager by his subordinate, which can provide an incentive for the manager to strategically hire a bad subordinate.

Whenever information is exchanged that is directly relevant to external observers beyond the outcome of negotiations, it may be desirable to keep negotiations secret even when no agents are involved. One example is central-bank policy making, where minutes of meetings would provide information to markets far beyond the information conveyed by the implemented policy. For a detailed discussion of central-bank secrecy, see Goodfriend (1986). Daughety and Reinganum (1995) argue that open pre-trial settlement negotiations would provide courts with information relevant to their eventual judgement, making settlement less likely. It is therefore efficient not to admit pre-trial negotiations as evidence, as is often the rule. Daughety and Reinganum (2000) show that if culpability of a defendant is correlated between two potential plaintiffs, one plaintiff's knowledge of a settlement between the other and the defendant can lead to inefficient litigation. Hence, it may sometimes be efficient to keep settlements secret.

9 Concluding remarks

The puzzle that we have attempted to solve in this paper is why agents who negotiate on behalf of principals often engage in inefficient "posturing" when their actions are observed. Our theory is based on the idea that agents differ in their ability to predict the opponent's bargaining position, and promote their reputation as skilled negotiators through their bargaining behavior. To isolate this effect, 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' reputational 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, union leaders are more likely to engage in posturing than are managers because the latter are typically evaluated on many criteria other than their performance in labor negotiatons. Also, studies have established 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. Thus, to the extent that politicians are concerned with their reputation within their own party while lawyers in criminal cases care about their market reputation, closed-door committees in politics seem easier to justify than closed-door trials in court.

(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 talent by working harder, monitoring may be desirable.

(5) if it is relatively easy to determine whether a bargaining outcome is efficient. This is the case in our model, where trade at any price is efficient, 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, they also cannot determine whether the agents' efforts were adequate (or whether they colluded). In this case, monitoring may 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 can 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 situations 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, union and management delegates whose moves are monitored often engage in tacit communication, i.e. 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, pp. 336-7). Such responses are sometimes interpreted as an attempt of 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. Quite generally, 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.

References

Abreu, Dilip, and Faruk Gul (2000): Bargaining and Reputation, *Econometrica* **68** (1), 85-118

Aghion, Philippe and Jean Tirole (1998): Formal and Real Authority, *Journal of Political Economy* **105** (1), 1-29

Bester, Helmut, and Józef Sákovics (1998): Delegated Bargaining and Renegotiation, mimeo, Free University of Berlin

Burtraw, Dallas (1993): Bargaining With Noisy Delegation, *Rand Journal of Economics* **24** (1) 40-57

Cai, Hongbin (2000): Bargaining on Behalf of a Constituency, *Journal of Economic Theory*

Caillaud and Rey (1994): Strategic Aspects of Vertical Delegation, *European Economic Review* **39**, 421-431

Chatterjee, Kalyan; and Larry Samuelson (1987): Bargaining with Two-Sided Incomplete Information: An Infinite Horizon Model with Alternating Offers, *Review of Economic Studies* **54** (2), 175-92.

Chatterjee, Kalyan; and William Samuelson (1983): Bargaining Under Incomplete Information, *Operations Research* **31** (5) pp. 835-851

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

Cramton, Peter (1985): Sequential Bargaining Mechanisms, in: Roth, Alvin E., ed.. Game-Theoretic Models of Bargaining. Cambridge; New York and Sydney: Cambridge University Press, 149–79.

Crawford, Vincent P. (1982): A Theory of Disagreement in Bargaining, *Econometrica* **50** (3) 607-638

Crémer, Jacques (1995): Arm's Length Relationships, *Quarterly Journal of Economics*, **110**, p. 275-95

Daughety, Andrew F., and Jennifer F. and Reinganum (1995): Keeping Society in the Dark: On the Admissibility of Pre-trial Negotiations as Evidence in Court, *Rand Journal of Economics* **26** (2), p 203-221

Daughety, Andrew F., and Jennifer F. and Reinganum (2000): Is Silence Golden? Confidentiality and Correlated Culpability, working paper, Vanderbilt University

Fershtman, Chaim, and Ehud Kalai (1997): Unobserved Delegation, *International Economic Review* **38** (4) 763-774

Fisher and Ury (1981): *Getting to Yes*, revised ed. 1991

Friebel, Guido, and Michael Raith (2000): Abuse of Authority and Hierarchical Communication, manuscript, University of Chicago

Goodfriend, Marvin (1986): Monetary Mystique: Secrecy and Central Banking, *Journal of Monetary Economics* **17** 63-92

Holmström, Bengt (1999): Managerial Incentive Problems: a Dynamic Perspective, *Review of Economic Studies* **66** (1) 169-182

Jones, Stephen R.G. (1989): Have Your Lawyer Call My Layer - Bilateral Delegation in Bargaining Situations, *Journal of Economic Behavior and Organization* **11**, 159-174

Katz, Michael (1991): Game-Playing Agents: Unobservable Contracts as Precommitments, *RAND Journal of Economics* **22**, 307–328

Kennan, John, and Robert Wilson (1993): Bargaining With Private Information, *Journal of Economic Literature* **31**, 45-104

Kroszner, Randall S., and Thomas Stratmann (2000): Does Political Ambiguity Pay? Corporate Campaign Contributions and the Rewards to Legislator Reputation, NBER working paper 7475

Myerson, Roger, and Mark Satterthwaite (1983): Efficient Mechanisms for Bilateral Trading, *Journal of Economic Theory*, **29** (2), 265–81.

Nierenberg, Gerard I. (1968): *The Art of Negotiating*, Prentice-Hall

Perry, Motty, and Larry Samuelson (1995): Open- versus Closed-Door Negotiations, *RAND Journal of Economics* **25** (2), 348–59.

Pruitt, Dean G. (1981) *Negotiation Behavior*, Academic Press

Rubin, Jeffrey Z., and Bert R. Brown (1975): *The Social Psychology of Bargaining and Negotiation*, Academic Press

Schelling, Thomas C. (1960): *The Strategy of Conflict* (Harvard)

Schotter, Andrew, Wei Zheng and Blaine Snyder (2000): Bargaining Through Agents: an Experimental Study of Delegation and Commitment, *Games and Economic Behavior* **30**, 248-92

Segendorff, Björn (1998): Delegation and Threat in Bargaining, *Games and Economic Behavior* **23**, 266-283

von Sydow, Emily (1995): Clamour for Openness, *European Voice* 9-15 November

Walton, Richard E.; and Robert B.McKersie (1965 [1991]): *A Behavioral Theory of Labor Negotiations*, (ILR Press)

A Proofs

Proof of Proposition 1: Because of the symmetry of the model, it suffices to consider the seller's side in determining equilibria, and we will drop the subscript s whenever the context is clear.

Open-door equilibrium: 1. Since the agents' bids are cheap talk, any pooling strategy with $a^t = a^b = a^w = a$ can be part of an equilibrium, supported by beliefs that treat each event as uninformative, i.e. $\widehat{q}(e) = q$ for all $e \in E^o$ that occur with positive probability, and $\widehat{q}(e) \leq q$ for events that occur with zero probability. These equilibria clearly violate the Consistency condition.

2. Now assume that at least two of the types t, d, w use different strategies. An agent of a weak seller can bid H, which results in the event HY or HN depending on the buyer's response, or he can bid L, which is always accepted by the buyer, i.e. leads to the event LY. Since the buyer accepts any acceptable price, partial separation of types implies that $\widehat{q}(HY)$, $\widehat{q}(HN)$ and $\widehat{q}(LY)$ cannot all be the same.

Suppose the ranking of beliefs is $\widehat{q}(HY) > \widehat{q}(LY) > \widehat{q}(HN)$. Inspection of Table 1 then shows that a type-w agent bids H, while a type-t agent bids L. A bad agent bids H if $p\widehat{q}(HY) + (1-p)\widehat{q}(HN) > \widehat{q}(LY)$, L if the reverse is true, and possibly mixes if the two payoffs are equal.

Consider the principal's beliefs when the agent plays as described above. Since B is weak with probability p and a good A_s knows B's type, we have $\Pr(HY|\text{good}) = p$, whereas $\Pr(HY|\text{bad}) = pa_s$ because HY is reached only if B is weak and A_s bids H, which occur independently. Applying (2), we have

$$\widehat{q}(HY) = \frac{pq}{pq + pa_s^b(1-q)} = \frac{q}{q + a_s^b(1-q)} > q.$$

Similarly, $\Pr(LY|\text{good}) = 1 - p$ and $\Pr(LY|\text{bad}) = 1 - a_s^b$ (a low price is always accepted by A_b), and therefore

$$\widehat{q}(LY) = \frac{(1-p)q}{(1-p)q + (1-a_s^b)(1-q)}. \tag{10}$$

Finally, since a good agent bids tough if and only if the rival is weak, a seller who observes that her agent's high price is rejected (event HN) can infer that her agent must be bad with certainty. Formally, we have $\Pr(HN|\text{good}) = 0$ and hence $\widehat{q}(HN) = 0$. A bad A_s's expected payoffs from bidding H or L are therefore given by

$$\begin{aligned}\pi_s^o(d,H,\mathbf{a}) &= p\widehat{q}(HY)+(1-p)\widehat{q}(HN) &&= p\frac{q}{q+a_s^b(1-q)} \quad \text{and} \\ \pi_s^o(d,L,\mathbf{a}) &= \widehat{q}(LY) &&= \frac{(1-p)q}{(1-p)q+(1-a_s^b)(1-q)}.\end{aligned} \tag{11}$$

A bad A_s will play H with positive probability only if $\pi_s^o(d,H,\mathbf{a}) \geq \pi_s^o(d,L,\mathbf{a})$, which using (11) is equivalent to

$$a_s^b \leq \frac{p-q+(1-p)pq}{1-q} < p. \tag{12}$$

For values of a^b that satisfy (12) and in particular $a^b < p \Leftrightarrow 1-a^b > 1-p$, it follows from (10) that $\widehat{q}(LY) < q$. We thus have $\widehat{q}(HY) > q > \widehat{q}(LY) > \widehat{q}(HN)$, as was to be confirmed.

Since a^b is bounded by p, A_s either mixes or always plays weak. He plays weak if (12) is violated even for $a_s^b = 0$, or equivalently, if (4) is violated. If $a^b \in (0,p)$ then (12) holds with equality and hence provides its exact value. Because of symmetry, the same conclusions hold for A_b's strategy.

This equilibrium satisfies Consistency: if B is weak, a good agent bids H, while a bad agent trades at H only with probability a_s^b, and otherwise at L. If B is tough, a good agent always trades at L, while a bad agent trades at L with probability $1-a_s^b$ and otherwise fails to trade. Hence, a good agent always attains a higher expected payoff. An analogous argument applies to the buyer.

3. The preceding arguments can be applied in the same way to show that an equilibrium exists in which $a^t = 1$, $a^w = 0$ and $a^b \in [0,1)$; and $\widehat{q}(HN) > q > \widehat{q}(LY) > \widehat{q}(HY)$. However, a similar argument as above shows that this equilibrium violates Consistency.

4. There exists no equilibrium in which the beliefs are ranked as in the above two, but with weak inequalities. Suppose, for example, that $\widehat{q}(HY) = \widehat{q}(LY) > \widehat{q}(HN)$. In this case, a type w agent would be indifferent between H and L, while types t and d would strictly prefer L. But then either all agents play L, leading to a pooling equilibrium, or a type w agent plays H with positive probability, in which case $\widehat{q}(HY) = 1$ and hence $\widehat{q}(HY)$ and $\widehat{q}(LY)$ cannot be equal. Apply similar arguments to eliminate other cases of weak inequalities.

5. Finally, no rankings other than $\widehat{q}(HY) > \widehat{q}(LY) > \widehat{q}(HN)$ or $\widehat{q}(HY) < \widehat{q}(LY) < \widehat{q}(HN)$ can occur in a separating equilibrium. If, for example, the beliefs were $\widehat{q}(HY) >$

$\widehat{q}(HN) > \widehat{q}(LY)$, then playing H would be a dominant strategy for all players, leading to a pooling equilibrium. Rankings involving weak inequalities can be ruled out with either this or the preceding kind of argument. Put differently, since playing L leads to LY for all types of agents, LY can in equilibrium never be the best or the worst outcome. ■

Closed-door equilibrium: The same arguments as above also apply to the closed-door case and lead to the conclusion that in any equilibrium, the agents must either use a pooling strategy or separate in one of two ways described. One exception is that with closed-door bargaining, an equilibrium in which no trade is the best outcome for both agents cannot exist, for any agent could reach this event simply by rejecting any offer his rival makes. At any rate, only an equilibrium with $\widehat{q}(H) > \widehat{q}(L) > \widehat{q}(\emptyset)$ and $a^w = 1, a^t = 0$ satisfies the Consistency criterion. We now show the existence and uniqueness of such an equilibrium. Its derivation is more complicated than in the closed-doors case because now, an agent's optimal strategy depends on the one chosen by the other agent.

To derive the equilibrium strategy of a bad agent, we calculate the beliefs $\widehat{q}(e)$, using the abbreviations a_s and a_d for a_s^b and a_b^b. A good A_s can arrive at event H only if B is weak. Conditional on this, a good A_s will bid H with certainty if he is to make an offer, while if A_b is to make an offer, only a bad A_b will bid H (with probability a_b), since a good A_b who knows that S is weak would bid L. Thus we have $\Pr(H|\text{good}) = p[1 + (1-q)(1-a_b)]/2$, and we similarly obtain $\Pr(H|\text{bad}) = p[a_s + (1-q)(1-a_b)]/2$. Hence

$$\widehat{q}(H) = \frac{[1 + (1-q)(1-a_b)]q}{[1 + (1-q)(1-a_b)]q + [a_s + (1-q)(1-a_b)](1-q)} = \frac{[1 + (1-q)(1-a_b)]q}{1 + (a_s - a_b)(1-q)}. \tag{13}$$

A good A_s trades at L if (1) A_s makes an offer and B is tough (probability $1-p$), or if (2) A_b makes an offer and either (i) B is tough (probability $1-p$) or (ii) B is weak but A_b is good and hence knows that S is weak (probability pq), or (iii) B is weak and A_b is bad but plays tough (probability $p(1-q)a_b$). Thus we have $\Pr(HL|\text{good}) = [1 - p + 1 - p + pq + p(1-q)a_b]/2$, and similarly obtain $\Pr(L|\text{bad}) = [1 - a_s + 1 - p + pq + p(1-q)a_b]/2$. Hence

$$\begin{aligned}\widehat{q}(L) &= \frac{[2(1-p) + pq + p(1-q)a_b]q}{[2(1-p) + pq + p(1-q)a_b]q + [2 - p - a_s + pq + p(1-q)a_b](1-q)} \\ &= \frac{[2(1-p) + pq + p(1-q)a_b]q}{2 - p + p(1-q)a_b - (1-q)a_s}. \end{aligned} \tag{14}$$

If a weak S observes no trade ($\emptyset$), she can infer that her own agent made an offer, since he would have accepted any offer by the buyer. But since a good agent always trades, her agent must be bad, i.e. $\widehat{q}(\emptyset) = 0$, as with open doors.

A bad A_s's expected payoffs from bidding H or L are then given by

$$\begin{aligned} \pi_s^c(d, H, \mathbf{a}) &= p\widehat{q}(H) + (1-p)\widehat{q}(\emptyset) = p\frac{[1 + (1-q)(1-a_b)]q}{1 + (a_s - a_b)(1-q)} \quad \text{and} \\ \pi_s^c(d, L, \mathbf{a}) &= \widehat{q}(L) = \frac{[2(1-p) + pq + p(1-q)a_b]q}{2 - p + p(1-q)a_b - (1-q)a_s}, \end{aligned} \tag{15}$$

and A_b's payoffs are obtained analogously. A_s bids H if $p\widehat{q}(H) > \widehat{q}(L)$ or if

$$\Delta(a_s, a_b) := p\widehat{q}(H) - \widehat{q}(L) > 0,$$

he bids L if $\Delta(a_s, a_b) < 0$, and possibly mixes if $\Delta(a_s, a_b) = 0$. Examining (13) and (14), it is straightforward to show that $\widehat{q}(H)$ is decreasing in a_s and $\widehat{q}(L)$ is increasing in a_s. Hence $\Delta(a_s, a_b)$ is decreasing in a_s. Moreover, $\Delta(a_s, a_b)$ is negative for $a_s = 1$:

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

A_s's best response to a_b therefore is L if $p\widehat{q}(H) < \widehat{q}(L) \Leftrightarrow \Delta(a_s, a_b) < 0$ even for $a_s = 0$ (where Δ is largest). Otherwise, we have $\Delta(0, a_b) > 0$, and monotonicity implies that there exists a unique $a_s^*(a_b) \in (0,1)$ defined by $\Delta(a_s^*(a_b), a_b) = 0$ or

$$a_s^*(a_b) = \frac{p(3-p)(2-q) - 2}{2(1-q)} + \frac{1}{2}a_b(1-p)[2 - p(3 - q - a_b(1-q))] \tag{16}$$

in the interior of the unit interval. Put differently, A_s's best response to a_b is given by $a_s^*(a_b)$ in (16) if it is positive, and zero otherwise.

Differentiate (16) with respect to a_b to obtain the slope of $a_s^*(a_b)$:

$$\frac{\partial a_s^*(a_b)}{\partial a_b} = \frac{1}{2}(1-p)[2 - p(3 - q - 2a_b(1-q))] \tag{17}$$

This slope is increasing in a_b; hence the minimal and maximal values are attained at $a_b = 0$ and $a_b = 1$, respectively. We then have

$$\begin{aligned} \frac{\partial a_s^*(1)}{\partial a_b} &= (1-p)\left[1 - \frac{1}{2}p(1+q)\right] < 1 \quad \text{and} \\ \frac{\partial a_s^*(0)}{\partial a_b} &= -p(1-p)\frac{3-q}{2} + (1-p) > -1, \end{aligned}$$

which implies that the slope of $a_s^*(a_b)$ is less than 1 in absolute terms. This means that there exists no more than one equilibrium with $a_s, a_b \in (0,1)$. Because the model is symmetric, uniqueness also implies that an interior equilibrium, if it exists, must be symmetric.

Next, $a_s = a_b = 0$ is an equilibrium if

$$a_s^*(0) = \frac{p(3-p)(2-q)-2}{2(1-q)} < 0,$$

the reverse of which is condition (6) stated in the proposition. Moreover, if (6) holds, then a symmetric interior equilibrium exists: Brouwer's fixpoint theorem implies that the function $\max\{0, \min\{a_s^*(a_b), 1\}\}$ must have a fixpoint on the unit interval, and since $a_s^*(0) > 0$ and $a_s^*(1) = [p(2+q) - 2q - p^2 q]/[2(1-q)] < 1$, the fixpoint must lie in the interior of $(0,1)$.

Finally, no asymmetric equilibrium with $a_s = 0$, $a_b > 0$ exists. Suppose it did. First, this would require (6) to hold. Second, the symmetric counterpart with $a_b = 0$, $a_s > 0$ would be an equilibrium as well. Thus, the agent's reaction curves would have to intersect three times: at a symmetric point in the interior of $(0,1)^2$, and at two points with $a_s > 0, a_b < 0$ and $a_s < 0, a_b > 0$, respectively. But since the reaction curves are convex, $a_s^*(a_b)$ would have to have an absolute slope exceeding -1 at the symmetric equilibrium, which we ruled out above. Hence, such asymmetric equilibria cannot exist.

We have thus shown existence, uniqueness and symmetry of a closed-door equilibrium. If (6) holds, bad agents mix, and otherwise they play weak. The mixing probability a^b as stated in the Proposition is the symmetric solution to (16) in $(0,1)^2$. ■

Proof of Lemma 1: Since $\widehat{q}(H)$ is increasing in a_b and for $a_s < p$, $\widehat{q}(L)$ is decreasing in a_b, it suffices to show the stated relationships for the case $a_b = 1$. In this case, $\widehat{q}(H)$ coincides with $\widehat{q}(HY)$. Moreover,

$$\begin{aligned} \widehat{q}(L) - \widehat{q}(LY) &= \frac{(2-p)q}{2 - pq - a_s(1-q)} - \frac{(1-p)q}{(1-p)q + (1-a_s)(1-q)} \\ &= \frac{q(1-q)(p-a_s)}{[2 - pq - a_s(1-q)][(1-p)q + (1-a_s)(1-q)]} > 0 \end{aligned}$$

Finally, $\widehat{q}(\emptyset) = \widehat{q}(HN) = 0$ as explained in text. ■

Proof of Proposition 2: For any a_b, A_s's best response as given by (16) is smaller than

a^o according to (3), as can be seen be examining the difference:

$$\begin{aligned} a^o - a_s^c(a_b) &= \frac{p-q+(1-p)pq}{1-q} - \frac{p(3-p)(2-q)-2}{2(1-q)} - \frac{a_b}{2}(1-p)[2-p(3-q-a_b(1-q))] \\ &= \frac{1-p}{2(1-q)}\left[2(1-p)(1-q)(1-a_b) + pq + pa_b(1-a_b)(1-q)^2\right] > 0 \end{aligned}$$

In particular, this must then also hold for a^c, which is a best response to itself. Notice that in describing the intuition behind this result in the text, we use Proposition 1, which assumes that $a_s, a_b < p$. In the proof here, however $a^c < p$ is not assumed but instead follows from $a^c < a^o$ shown here, and $a^o < p$ shown earlier. ■

Proof of Lemma 2: By definition,

$$\widehat{q}(e) = \frac{\Pr(e, \text{good})}{\Pr(e) + \Pr(e, bad)},$$

or alternatively

$$[1 - \widehat{q}(e)] \Pr(e, \text{good}) = \widehat{q}(e) \Pr(e, bad)].$$

Using this relationship, we have

$$\begin{aligned} \text{Var}(\widehat{q}) &= \sum_e [\widehat{q}(e)]^2 \Pr(e, bad) + [\widehat{q}(e) - 1]^2 \Pr(e, \text{good}) \\ &= \sum_e [\widehat{q}(e)]^2 \Pr(e, d) + [1 - \widehat{q}(e)] \widehat{q}(e) \Pr(e, bad) \\ &= \sum_e \Pr(e, bad) \widehat{q}(e) \\ &= (1-q) \sum_e \Pr(e|\text{bad}) \widehat{q}(e). \end{aligned}$$

Proof of Proposition 3: The set of observable events is coarser with closed doors than with closed doors. In particular, we have $H = \{HY, YH\}$ and $L = \{LY, YL\}$.

Consider any events e_1 and e_2 for which e_2 occurs when A_b makes an offer, i.e. if $\Pr(e_2|\text{good}) = \Pr(e_2|\text{bad}) =: \Pr(e_2)$ and $\widehat{q}(e_2) = q$. We then have

$$\Pr(e|\text{bad}) \widehat{q}(e) > \Pr(e_1|\text{bad}) \widehat{q}(e_1) + \Pr(e_2|\text{bad}) \widehat{q}(e_2)$$

where $e = \{e_1, e_2\}$. The result then follows because of Lemma 2.

Under the stated assumptions about e_1 and e_2, we have $\Pr(e|\text{bad}) = \Pr(e_1|\text{bad}) + \Pr(e_2)$, $\widehat{q}(e_1) = \Pr(e_1|\text{good})q/[\Pr(e_1|\text{good})q + \Pr(e_1|\text{bad})(1-q)]$, and an analogous condition for $\widehat{q}$. Thus

$$\begin{aligned}\Pr(e|\text{bad})\widehat{q}(e) &= \frac{[\Pr(e_1|\text{bad}) + \Pr(e_2)][\Pr(e_1|\text{good}) + \Pr(e_2)]q}{[\Pr(e_1|\text{good}) + \Pr(e_2)]q + [\Pr(e_1|\text{bad}) + \Pr(e_2)](1-q)} \\ &= \frac{[\Pr(e_1|\text{bad}) + \Pr(e_2)][\Pr(e_1|\text{good}) + \Pr(e_2)]q}{\Pr(e_1|\text{good})q + \Pr(e_1|\text{bad})(1-q) + \Pr(e_2)}\end{aligned}$$

and

$$\Pr(e_1|\text{bad})\widehat{q}(e_1) + \Pr(e_2|\text{bad})\widehat{q}(e_2) = \frac{\Pr(e_1|\text{bad})\Pr(e_1|\text{good})q}{\Pr(e_1|\text{good})s + \Pr(e_1|\text{bad})(1-q)} + \Pr(e_2)q,$$

leading to

$$\begin{aligned}&\Pr(e|\text{bad})\widehat{q}(e) - [\Pr(e_1|\text{bad})\widehat{q}(e_1) + \Pr(e_2|\text{bad})\widehat{q}(e_2)] \\ =\ &\frac{q^2(1-q)[\Pr(e_1|\text{good}) - \Pr(e_1|\text{bad})]^2 \Pr(e_2)}{[\Pr(e_1|\text{good})q + \Pr(e_1|\text{bad})(1-q)][\Pr(e_1|\text{good})q + \Pr(e_1|\text{bad})(1-q) + \Pr(e_2)]} > 0\end{aligned}$$

■

Proof of Proposition 4: We prove this result by providing two examples:

1. If $p = 2/3$ and $q = 1/2$ we have $a^o = 0.56$ and $a^c = 0.36$ according to Proposition 1. The corresponding variance of the estimation error is 0.2317 for closed doors, using (8). Since a increases to 0.56 with open doors, the resulting variance in equilibrium is 0.2321 (using (7)), i.e. larger than in the closed-door equilibrium.

2. If $p = 2/3$ and $q = 1/4$, then we obtain $a^o = 0.63$ and $a^c = 0.52$. Here, evaluating (7) and (8) leads to $\mathrm{Var}(u^o) = 0.1804$ and $\mathrm{Var}(u^o) = 0.1803$. Here, open-door bargaining is (slightly) more informative than closed-door bargaining. ■

Proof of Proposition 5: If B can observe the agents' bids, the situation for B and A_b is as with open doors, hence $a_b = a^o$. For S, bargaining is with closed doors. We use the closed-door best response (16) to show that $a_s^*(a_b)$ is contained in (a^c, a^o) for $a_b = a^o$. Substitute $a_b = a^c$ as given in Proposition 1 into (17) to obtain

$$\frac{\mathrm{d}a_s^*(a^c)}{\mathrm{d}a_b} = 1 - \frac{1}{2}\sqrt{p}\sqrt{4(2-p) - p(1-p)(1-q)[3+q+p(1-q)]} \tag{18}$$

This slope is positive if

$$p\{4(2-p)-p(1-p)(1-q)[3+q+p(1-q)]\} = 4p(2-p)-p^2(1-p)(1-q)[3+q+p(1-q)] < 4,$$

which is the case because $p(2-p)$ in the first term does not exceed 1 and the second term is negative. Thus, $a_s^*(a_b)$ is increasing in a_b at a^c, and since $a_s^*(a_b)$ is convex in a_b (cf. the proof of Proposition 1), it follows that $a_s^*(a_b)$ must exceed a^c for any $a_b > a^c$, in particular $a_b = a^o$. Finally, Proposition 2 implies that $a_s^*(a_b) < a^o$. ∎

Proof of Proposition 6: Consider S's expected payoff ex ante, i.e. before she knows her own reservation value. If both S and B turn out to be tough, no trade is possible. If S is tough and B weak (which occurs with probability $p(1-p)$), then only trade at H is possible, leading to a payoff of $H-\bar{s}$ for S. Trade at H will always occur except if A_b is to make an offer, A_b is bad, and chooses to play tough (bid L). Similarly, if B is tough and S weak, trade at L will occur except if A_s is bad, makes the offer and plays tough. If both principals are weak (with probability p^2), then trade at H occurs if A_s makes an offer and is either good or plays tough, and if a bad A_b makes the offer and plays weak. The probability that trade at L occurs can be calculated similarly. The expected payoff for S therefore is

$$\begin{aligned}\Pi_s(\mathbf{a}) &= p(1-p)[1-\frac{1}{2}(1-q)a_b](H-\bar{s})+p(1-p)[1-\frac{1}{2}(1-q)a_s](L-\underline{s}) \\ &\quad +\frac{p^2}{2}\{[q+(1-q)a_s+(1-q)(1-a_b)](H-\underline{s}) \\ &\quad +[q+(1-q)a_b+(1-q)(1-a_s)](L-\underline{s})\}, \end{aligned} \tag{19}$$

and B's payoff can be calculated in similar fashion.

Then we have

$$\frac{\partial \Pi_s(\mathbf{a})}{\partial a_s} = p(1-q)[p(H-\underline{s})-(L-\underline{s})] \quad \text{and} \tag{20}$$

$$\frac{\partial \Pi_s(\mathbf{a})}{\partial a_b} = -p(1-q)[(1-p)(H-\bar{s})+p(H-\underline{s})-(L-\underline{s})]. \tag{21}$$

Since according to Proposition 5 both a_s and a_b exceed a^c, (20) and (21) imply that $p(H-\underline{s}) > L-\underline{s}$ is a necessary but not sufficient condition for unilateral monitoring to be profitable for S (it is easy to construct a counterexample where S is almost indifferent about her own agent's play but loses profit if a_b increases). If this condition holds and the reservation prices are symmetric, then we also have $\partial\Pi_b(\mathbf{a})/\partial a_b > 0$. If S monitors the agents, B would want to monitor them as well, since a shift from unilateral monitoring by S to open-door bargaining would lead to an increase in a_b, while a_s remains at a^o.

Finally, to prove the existence of parameters for which S gains by monitoring unilaterally, consider the example $\underline{s} = 0$, $L = 0.25$, $\underline{b} = \bar{s} = 0.5$, $H = 0.75$, $\bar{b} = 1$, $p = 0.5$ and $q = 0.25$. The mixing probabilities for open and closed doors are $a^o = 0.417$ and $a^c = 0.151$, respectively. Substituting these values into (19), we obtain an ex-ante expected profit of 0.461 with open doors and 0.486 with closed doors. Since $p(H - \underline{s}) - (L - \underline{s}) = 0.125$, S would want her own bad agent to be tough. If S monitors unilaterally, then $a_s = a^o = 0.417$ and the corresponding a_b is 0.206 according to (16). S's expected profit now is 0.49, i.e. higher than with closed doors, while B's expected profit is 0.451, i.e. lower than with open or closed doors. Hence, S would want to monitor the agents unilaterally, but B would then want to monitor them as well. ■

B An endogenization of feasible price bids

In this appendix, we discuss in more detail our assumption that agents of tough principals must play tough, whereas agents of weak principals are given discretion over what prices they suggest or accept. This assumption is an important ingredient of our theory.

In our model, statistical separation between good and bad agents occurs not per se because good agents know the rival principal's bargaining position and bad agents do not, but because good agents are better able to predict the rival *agent's* bargaining behavior. To see this, suppose the buyer's agent was instructed always to bid L. Then a good and a bad seller agent would have the same information about the probabilities of the different events. But then the payoff functions of good and bad agents would be identical, leading to a pooling equilibrium. Thus, in our theory, knowledge about the rival principal is relevant because it allows a good agent to better predict the rival agent's actions. This requires that agents of weak and of tough (rival) principals behave differently, a minimal condition for which is that agents of weak and of tough principals choose from different sets of feasible prices.

If it is prohibitively costly to a principal not to ratify an agreement negotiated by the agents, then each principal will naturally want to restrict the set of prices which her agent is allowed to suggest or accept to those prices at which trade is profitable for the principal. The question, however, is whether a weak principal would indeed want to give her agent discretion over the choice of prices.

If the principal and her agent had the same knowledge about the rival, there would be no reason to give the agent discretion (or even delegate bargaining in the first place); the principal would simply require her agent to to always bargain tough or weak, depending on her expectations about the rival's bargaining position (the relevant condition is given by equation (9)). Thus, the only reason for delegation is that good agents can on average achieve better outcomes if they use their knowledge about the rival. On the other hand, since a principal doesn't know her agent's type, the cost of delegation, i.e. of giving discretion, is that a bad agent's bargaining behavior is suboptimal from the principal's point of view. It follows that that delegation is optimal if the share of good agents q exceeds some lower bound. We now derive this lower bound formally.

We generalize our basic model in the following way: suppose at stage 4 of the timing described in Section 2, each principal independently chooses the set of prices her agent is allowed to agree to. This set depends on the principal's reservation price. A tough seller will restrict her agent's feasible prices to the set $\{H\}$, whereas a weak seller must choose one of the sets $\{H\}$, $\{L\}$ and $\{H, L\}$. Consider any general strategy (a_b^d, a_b^w) that A_b might use, depending on whether he is bad or whether he knows that S is weak. Let $h = H - \underline{s}$ and $l = L - \underline{s}$. If a weak S instructs her agent to always play H, then her expected payoff is

$$\frac{p}{2}\{h + q[a_b^w l + (1 - a_b^w)h] + (1 - q)[a_b^b l + (1 - a_b^b)h]\} + (1 - p)\frac{l}{2}. \tag{22}$$

Similarly, if she instructs her agent to always play L, her expected payoff is

$$\frac{p}{2}\{l + q[a_b^w l + (1 - a_b^w)h] + (1 - q)[a_b^b l + (1 - a_b^b)h]\} + (1 - p)l. \tag{23}$$

Finally, if S gives A_s discretion, then her agent plays a pure strategy according to Proposition 1 if he is good, and plays H with some probability a_s if he is bad. The resulting expected payoff for S is

$$\frac{p}{2}\{qh + (1 - q)[a_s h + (1 - a_s)l] + q[a_b^w l + (1 - a_b^w)h] + (1 - q)[a_b^b l + (1 - a_b^b)h]\}$$
$$+\frac{l - p}{2}(1 + q + (1 - q)(1 - a_s))l. \tag{24}$$

Given the one-shot nature of our game, it turns out that A_b's strategy is irrelevant for S's choice of price set for her own agent; that is, the differences between (22), (23) and (24) do not depend on a_b^w or a_b^b.

S prefers that her agent always bid H rather than L if the difference between (23) and (22) is positive. This difference has the same sign as $ph > l$, which is equivalent to (9). Suppose this condition holds. Then delegation is optimal if and only if the difference between (24) and (23), which is

$$\frac{1}{2}\{pq(h - l) - (ph - l)[1 - (1 - q)a_s]\}, \tag{25}$$

is positive. Since (25) is increasing in a_s, a sufficient condition for (25) to be positive is

$$pq(h - l) > ph - l, \tag{26}$$

which implies a lower bound for q. If, on the other hand, $ph < l$, then S prefers that her agent always bid L rather than H. In this case, delegation is optimal if and only if the difference between (24) and (22), or

$$\frac{1}{2}[pq(H - L) - as(1 - q)(L - Hp)], \tag{27}$$

is positive. Since (27) is decreasing in a_s, a sufficient condition for (27) to be positive is

$$pq(h - l) > (1 - q)(ph - l), \tag{28}$$

which again implies a lower bound for q. Since (26) implies (28), a sufficient condition for optimal delegation that does not depend on q or a_s is (26), or $q \geq (ph - l)/[p(h - l)]$.

C Two-stage game: details

Consider a twice-repeated version of the bargaining game of Section 2. That is, an agent is chosen to make an offer, the other accepts or rejects, and if he rejects, a second round is played in which again one agent is randomly chosen to make an offer. With open doors, principals observe the entire sequence of agents' actions. We denote by, say, HNYL the event where in the first round, A_s bids H, which A_b rejects, whereas in the second round A_b gets to make an offer and bids L, which A_s accepts; and describe other possible sequences of bids in similar fashion. With closed doors, principals observe only the eventual outcome (H, L or $\emptyset$) along with the round in which it is reached (1 or 2), and we denote these events by H2 etc. Table 2 lists all closed-door outcomes along with the sequences of bids that lead to these outcomes, i.e. the corresponding open-door events.

Table 2: Closed-door and corresponding open-door events

Closed doors	Open doors
1H	LY, YL
1L	HY, YL
2H	HNHY, HNYH, NLHY, NLYH
2L	HNLY, HNYL, NLLY, NLYL
$\emptyset$	HNHN, NLHN

Just as in the one-stage game, the possible separating equilibria must entail beliefs that, at each stage of the game (except second-round responses), force a bad agent to choose between a sure option and another that leads to either a better or a worse reputation than the sure option (see the proof of Proposition 1). This will induce good agents to choose pure strategies that depend on their information about the rival. As above, we focus on a weak seller's perspective and apply the Consistency requirement, which means that in equilibrium, the seller's reputation is generally increasing in the price at which he trades.

A good A_s who knows that B is weak will then bid H in both rounds and reject a bid of L by A_b in the first round. A good A_s who knows that B is tough, on the other hand, will concede immediately, that is both offer and accept L in the first round. A bad agent may randomize at each point in the game, and his second-stage decisions may depend on the outcome of the first stage. A bad agent's strategy is therefore characterized by the vector $(a_1, a_{HN}, a_{NL}, a_n)$, where a_1 is the probability that A_s bids H in $t = 1$, a_n is the probability that he rejects A_b's bid of L in $t = 1$, a_{HN} is the probability that A_s bids H in $t = 2$ if the outcome is HN, and a_{NL} is defined similarly.

Let r_1 be the probability that the agent of a *weak* B bids L in $t = 1$ from a weak S's perspective who does not know the type of A_b. That is, if e.g. a good A_b always bids L (knowing that S is weak) and a bad A_b bids L with probability z_1, then $r_1 = q + (1 - q)z_1$. Similarly, let r_{HN} and r_{NL} be the probabilities that a weak A_b bids L in $t = 2$ after the outcomes HN and NL in $t = 1$, respectively; and let r_n be the probability that A_b rejects H in $t = 1$, all evaluated from S' perspective.

With this notation, we can then describe, for each observable event e, the probabilities that a good and a bad agent reach e, and the determine the posterior probability $\widehat{q}(e)$ with which S believes to have a good agent upon observing e. These probabilities are given in Table 3 for open-door bargaining, and the corresponding posterior beliefs $\widehat{q}(e)$ are computed using (2). For closed-door bargaining, first calculate the event probabilities for each agent type and each closed-door event by adding up the open-door probabilities that correspond to that event according to Table 2, and then compute the beliefs using (2) and the closed-door event probabilities.

Open-door bargaining: To determine the open-door equilibrium, consider A_s' decision whether to bid H or L in the second stage. As in the one-shot game, rejection of H by A_b is sure information that A_s is bad ($\widehat{q}(HNHN) = \widehat{q}(NLHN) = 0$), for a good agent never bids H unless he knows that B is weak and A_b accepts. On the other hand, giving in by bidding L, too, is sure information that A_s is bad ($\widehat{q}(HNLY) = \widehat{q}(NLLY) = 0$), because a good agent never lets bargaining proceed to the second round unless he knows that B is weak. For this reason, a bad A_s always chooses to bid H in $t = 2$, hoping that A_b will accept, i.e. we have $a_{HN} = a_{NL} = 1$. Thus, by bidding H or rejecting L in $t = 1$ a bad agent in effect signals to S that he believes B to be weak. He is then locked into this strategy and must then do whatever a good agent would do, even if A_b's tough stance in $t = 1$ is bad news for A_s.

Table 3: Open-door Event probabilities $\Pr(e|k)$ conditional on agent type

Event e	Good agent ($k = s$)	Bad agent ($k = d$)
LY	$\frac{1}{2}(1-p)$	$\frac{1}{2}(1-a_1)$
YL	$\frac{1}{2}(1-p)$	$\frac{1}{2}(1-a_n)(1-p+pr_1)$
HY	$\frac{1}{2}p(1-r_n)$	$\frac{1}{2}pa_1(1-r_n)$
YH	$\frac{1}{2}p(1-r_1)$	$\frac{1}{2}p(1-r_1)$
$HNLY$	0	$\frac{1}{2}a_1(1-p+pr_n)\frac{1}{2}(1-a_{HN})$
$HNYL$	$\frac{1}{4}pr_1r_{HN}$	$\frac{1}{4}a_1(1-p+pr_nr_{HN})$
$NLLY$	0	$\frac{1}{4}a_n(1-p+pr_1)(1-a_{NL})$
$NLYL$	$\frac{1}{4}pr_1r_{NL}$	$\frac{1}{4}a_n(1-p+pr_1r_{NL})$
$HNHY$	$\frac{1}{4}pr_n$	$\frac{1}{4}pa_1r_na_{HN}$
$HNYH$	$\frac{1}{4}pr_n(1-r_{HN})$	$\frac{1}{4}pa_1r_n(1-r_{HN})$
$NLHY$	$\frac{1}{4}pr_1$	$\frac{1}{4}pr_1a_na_{NL}$
$NLYH$	$\frac{1}{4}pr_1(1-r_{NL})$	$\frac{1}{4}pr_1a_n(1-r_{NL})$
$HNHN$	0	$\frac{1}{4}(1-p)a_1a_{HN}$
$NLHN$	0	$\frac{1}{4}(1-p)a_na_{NL}$

In $t=1$, a bad A_s mixes between H and L if his expected payoffs for each action are the same. This indifference is expressed by

$$\begin{aligned} & (1-p)\frac{\widehat{q}(HNHN)+\widehat{q}(HNYL)}{2}+ \\ & p\left\{(1-r_1)\widehat{q}(HY)+\frac{r_1}{2}[\widehat{q}(HNHY)+r_{HN}\widehat{q}(HNYL)+(1-r_{HN})\widehat{q}(HNYH)]\right\} \\ = \; & \widehat{q}(LY). \end{aligned} \tag{29}$$

The first term on the l.h.s. corresponds to the case that B is tough, in which case A_b always rejects H. In the second round, bargaining results in disagreement (HNHN) if A_s makes an offer (because he must bid H, cf. the discussion above), or in trade at L (HNYL) if A_b makes an offer. The second term corresponds to the case that B is weak, in which case A_b accepts H with probability $1-r_n$ and otherwise rejects. The second-period expected payoff (term in [ ]) is calculated in similar fashion, where again A_b randomizes (from S' perspective) between H and L.

A_s mixes between accepting and rejecting a price of L in $t=1$ if this leads to the same expected payoffs, which are calculated in the same way as in (29):

$$(1-p)\frac{\widehat{q}(NLHN)+\widehat{q}(NLYL)}{2}+\frac{p}{2}[\widehat{q}(NLHY)+r_{NL}\widehat{q}(NLYL)+(1-r_{NL})\widehat{q}(NLYH)]=\widehat{q}(YL). \tag{30}$$

In a symmetric equilibrium, a bad A_b bids L in $t=1$ with probability a_1, bids L in $t=2$ with probability a_{HN} if the first-period outcome is LN, bids L in $t=2$ with probability a_{NL} if the first-period outcome is HY, and rejects H in $t=1$ with probability a_n. Taking expectations (from S's perspective) over the types of A_b, we then have:

$$r_1=q+(1-q)a_1, \quad r_{HN}=q+(1-q)a_{NL}, \quad r_{NL}=q+(1-q)a_{HN}, \quad r_n=q+(1-q)a_n. \tag{31}$$

An (interior) symmetric open-door equilibrium is then determined by substituting (31) (with $a_{HN}=a_{NL}=1$) into (29) and (30) and solving these for a_1 and a_n. The solution can be determined only numerically. For e.g. $p=3/5$ and $q=1/2$, we obtain a unique interior solution at $a_1=0.53$ and $a_n=0.47$. The corresponding (non-zero) beliefs according to Table 5 are

$$\begin{gathered} \widehat{q}(HY)=0.65, \quad \widehat{q}(LY)=0.46, \quad \widehat{q}(YL)=0.47 \\ \widehat{q}(HNHY)=0.65, \quad \widehat{q}(HNYL)=0.499, \quad \widehat{q}(NLHY)=0.68, \quad \widehat{q}(NLYL)=0.53 \end{gathered} \tag{32}$$

It is straightforward to verify that given these beliefs, good agents indeed choose the pure strategies postulated above. Notice that trading at H in $t = 2$ leads to a reputation no worse (or even better) than trading at H in $t = 1$. However, since even a weak B's agent might bid L in $t = 2$, a good A_s knows B is weak always prefers to trade at H in 1. Notice also that accepting L in $t = 2$ is quite uninformative about A_s: on one hand, by delaying bargaining, A_s signaled his belief that B is weak. On the other hand, since A_b's bid of L is accepted by any agent, there is not much information to judge A_s by.

Closed-door bargaining: with closed doors, the fact that any offer made by A_b in $t = 2$ is accepted has two consequences: first, there is no longer a lock-in effect in the second round: since S cannot see who makes an offer, bidding L is now an option for a bad A_s who bid H or rejected L in $t = 1$. Second, no trade is again sure information that A_b is bad, because disagreement means that A_s must have made an offer in $t = 2$. Hence, bidding H can lead to a positive reputation for A_s only if the bid succeeds, with requires that B is weak.

If the outcome of the first round is HN, then in $t = 2$, A_s' subjective probability that B is weak reduces to

$$p_{HN} = \frac{pr_n}{pr_n + 1 - p},$$

reflecting the fact that A_b always rejects A_s's bid of H if B is tough but only with probability r_n if B is weak. A bad A_s is indifferent between H and L in $t = 2$ if

$$p_{HN}\widehat{q}(2H) = \widehat{q}(2L) \tag{33}$$

Similarly, after NL in $t = 1$, A_s believes that B is weak with probability

$$p_L = \frac{pr_1}{pr_1 + 1 - p},$$

and is indifferent between H and L in $t = 2$ if

$$p_L\widehat{q}(2H) = \widehat{q}(2L) \tag{34}$$

In $t = 1$, A_s is indifferent between H and L if

$$\begin{aligned}\widehat{q}(1L) \quad = \quad & \frac{1-p}{2}[a_{HN}\widehat{q}(\emptyset) + (1 - a_{HN})\widehat{q}(2L) + \widehat{q}(2L)] + \frac{p}{2}\{(1 - r_n)\widehat{q}(1H) \\ & + r_n[a_{HN}\widehat{q}(2H) + (1 - a_{HN})\widehat{q}(2L) + r_{HN}\widehat{q}(2L) + (1 - r_{HN})\widehat{q}(2H)]\}\end{aligned} \tag{35}$$

This condition is constructed similarly as (29), except that now a_{HN} and a_{NL} are not necessarily equal to 1. Finally, if A_b bids L in $t = 1$, then A_s' subjective assessment that B is weak is p_L as defined above, and A_s is indifferent between rejecting and accepting if

$$\begin{aligned} & \frac{1-p_L}{2}[a_{NL}\widehat{q}(\emptyset) + (1-a_{NL})\widehat{q}(2L) + \widehat{q}(2L)] + \\ & \frac{p_L}{2}[a_{HN}\widehat{q}(2H) + (1-a_{HN})\widehat{q}(2L) + r_{HN}\widehat{q}(2L) + (1-r_{HN})\widehat{q}(2H)] \\ = \; & \widehat{q}(1L). \end{aligned} \tag{36}$$

Numerical solution of (33) through (36), substituting (31), leads to the unique interior solution

$$a_1 = 0.40, \quad a_n = 0.40, \quad a_{HN} = 0.79, \quad a_{NL} = 0.22 \tag{37}$$

and the corresponding posterior beliefs are

$$\widehat{q}(1H) = 0.59, \quad \widehat{q}(1L) = 0.42, \quad \widehat{q}(2H) = 0.81, \quad \widehat{q}(2L) = 0.41, \quad \widehat{q}(\emptyset) = 0 \tag{38}$$

A direct comparison of the beliefs in (32) and (38) is not very insightful given the change in A_s's equilibrium strategy. More relevant for understanding the change in bargaining behavior is to evaluate the closed-door beliefs using the equilibrium strategies for open-door bargaining. These are

$$\widehat{q}(1H) = 0.57, \quad \widehat{q}(1L) = 0.46, \quad \widehat{q}(2H) = 0.68, \quad \widehat{q}(2L) = 0.51, \quad \widehat{q}(\emptyset) = 0 \tag{39}$$

Compare these with (32): For given open-door strategies, trading at H in $t = 2$ is as valuable with closed as with open doors, (compare $\widehat{q}(..HY)$ and $\widehat{q}(..YH)$ with $\widehat{q}(2H)$). But with $a_{HN} = a_{NL} = 1$, $\widehat{q}(2L)$ is now simply the mean of $\widehat{q}(NLYL)$ and $\widehat{q}(HNYL)$, which makes bidding in L in $t = 2$ attractive. In $t = 1$, trading at L has much the same effect as with open doors (compare $\widehat{q}(LY)$ and $\widehat{q}(YL)$ with $\widehat{q}(1L)$). But bidding H is less attractive, because $\widehat{q}(1H)$ is now a mean of $\widehat{q}(HY) > q$ and $\widehat{q}(YH) = q$. Hence, in both rounds, a bad A_s has an incentive to bargain less aggressively with closed than with open doors.

D Continuous types and prices: details

Consider the following continuous version of the model of Section 2. The seller's reservation price s, the buyer's reservation price b are both uniformly distributed on the open unit interval. Bargaining is conducted in a one-shot game in which one party is chosen to make a price bid, and the other accepts or rejects. The seller's agent A_s can bid any price in the interval $[s, 1]$ and accept any price no less than s. The buyers's agent A_b can bid any price in the interval $[0, b]$ and accept any price no greater than b. There are good and bad agents, just like in the basic model. As usual, we focus on the seller side in all of our discussion.

The analog of the pure strategies in the basic model is that a good seller agent will simply bid the buyers reservation price if trade at that price is profitable for the seller. However, sometimes it is not (if $b < s$); hence it is clear that even a good agent sometimes fails to reach an agreement. If a good agent knows that trade is not possible, it seems reasonable to assume that he should be able to distinguish himself from a bad agent by signaling his information directly instead of demanding some price that is rejected by the buyer. Hence, as an alternative to bidding a price, we allow each agent to announce that trade is not possible (we denote this outcome by $\emptyset$) and thus terminate bargaining immediately.[28] If a good agent bids the buyer's reservation price if trade is possible and announces $\emptyset$ if not, then the seller can infer that her agent must surely be bad if his price is rejected by the buyer.

With these assumptions, a bad agent's strategy is then characterized by a probability ν that he announces $\emptyset$, and a distribution over feasible prices in case he does bid a price (with probability $1 - \nu$). From the seller's point of view (who doesn't know the buyer's reservation price), the distribution of prices a good agent would bid is uniform between s and 1, i.e. does not have any mass points. It follows that the distribution of a bad agent's prices cannot have mass points either, for the observation of such price bids would be certain information that the agent is bad. Hence, a bad agent's price distribution is continuous over some range of prices and can be described by the density $f(p, s)$ which depends on the seller's reservation price.

A good agent announces $\emptyset$ if $b < s$, i.e. with probability s, whereas a bad announces $\emptyset$ with probability ν. The seller's posterior belief about A_s upon observing $\emptyset$ therefore is

$$\widehat{q}(\emptyset) = \frac{qs}{qs + (1 - q)\nu} \tag{40}$$

[28] The same results are obtained if good agents who know that trade is not possible simply bid their principal's reservation price.

A good agent announces some price $p > s$ with density 1, corresponding to the density of the buyer's valuations. A bad agent announces a price with probability $1 - \nu$, and conditional on that, with density $f(p, s)$. Hence, the seller's posterior belief about A_s upon observing trade at price p is

$$\widehat{q}(p) = \frac{q}{q + (1-q)(1-\nu)f(p,s)} \tag{41}$$

(recall that a rejected price always means $\widehat{q} = 0$). All prices p over which a bad agent mixes must lead to the same expected payoff $(1-p)\widehat{q}(p)$. In equilibrium, ν, which leads to a payoff of $\widehat{q}(\emptyset)$, must also be positive.[29] Hence, we require that

$$(1-p)\widehat{q}(p) = \widehat{q}(\emptyset) \tag{42}$$

for all p, which after substituting (40) and (41) leads to

$$f(p,s) = (1-s)\frac{\nu(1-p)(1-q) - pqs}{s(1-q)(1-\nu)}, \tag{43}$$

wherever this expression lies within the unit interval. Thus, $f(p, s)$ is linear and decreasing in p. Since (43) is negative for $p = 1$, the maximal price a bad agent bids is always less than 1. Denote this maximal price by $p_{max}(s, \nu)$. Next, ν can be determined as the solution to

$$\int_s^{p_{max}(s,\nu)} f(p,s,\nu)dp = 1,$$

which leads to

$$\nu = s\frac{1 - q - qs^2 + \sqrt{1 - q^2(1+s^2)}}{(1-q)(1+s^2)} \tag{44}$$

Substitute (44) into (43) to obtain the price density f as a function of p and s only:

$$f(p,s) = \frac{1 - p - q(1+s^2) + (1-p)\sqrt{1-q^2(1+s^2)}}{1 - s + s^2 - q(1-s)(1+s^2) - u\sqrt{1-q^2(1+s^2)}}. \tag{45}$$

Finally, substitute (44) and (45) into (40) and (41) to obtain

$$\widehat{q}(\emptyset) = \frac{q(1+s^2)}{1+\sqrt{1-q^2(1+s^2)}} \quad \text{and} \quad \widehat{q}(p) = \frac{q(1+s^2)}{(1-s)(1+\sqrt{1-q^2(1+s^2)})}. \tag{46}$$

[29] If $\nu = 0$, only good agents would announce $\emptyset$. But then $\widehat{q}(\emptyset) = 1$, inducing bad agents to announce $\emptyset$. Similarly, if $\nu = 1$, then trade at any price would lead to $\widehat{q} = 1$, inducing agents to announce some price.

It is straightforward to show that $\widehat{q}(\emptyset)$ is increasing in s, and that $\widehat{q}(p)$ is increasing in p. Moreover, $\widehat{q}(s)$ is always greater than $\widehat{q}(\emptyset)$, which implies that trading at a profit for the principal always leads to a better reputation than announcing that trade is not possible.

Unfortunately, we cannot obtain a closed-door equilibrium for this model. Recall that determining an expression for f from (40), (41) and (42) above was possible because neither $\widehat{q}(\emptyset)$ nor $\widehat{q}(p)$ directly depend on p. With closed doors, $\widehat{q}(\emptyset)$ and $\widehat{q}(p)$ also depend on the expected behavior of the buyer's agent. Forming these expectations requires integrating the buyer's price density function over both prices and the buyer's valuations. But while f above is linear in p, it is highly nonlinear in s; and since the densities in the closed-door case are unlikely to have a simpler structure, there is little hope of obtaining a closed-form solution for f.

E Simultaneous bids: overview of the results

Consider a one-shot bargaining game in which the agents simultaneously submit bids. If the price bids are equal (H or L), the good is traded at this price. If both agents submit weak bids (A_s bids L and A_b bids H), the good is traded at one of the prices with equal probability. If both agents play tough, no trade takes place. With open doors, the principals can see the actual price bids, while with closed doors they can see only the outcome, trade at some price or disagreement.

This model leads to the same main result that bargaining behind closed doors is more efficient. As in the sequential model, good agents bid tough or weak depending on whether the rival is weak or tough. Consequently, if both agents play weak (which can occur only if both principals are weak), then with open doors both agents are revealed to be bad because each misjudged the rival. With closed doors, the principals observe only trade at L or H in this case. Hence, for a bad agent of a weak seller, selling at H is less of a good signal with closed than with open doors because trade at H can also occur if both agents make weak bids. Bidding L is more attractive with closed doors because the risk of being revealed as bad is absent. Finally, no trade is equally informative with closed as with open doors, since the principals can infer that A_s must have bid H and A_b L even if they cannot see the bids.[30] Thus, while some details differ from the sequential game, the intuition for the main result is the same: with closed doors, successfully bidding a high price is less attractive than with open doors, unsuccessfully bidding a high price is equally attractive, and bidding a low price is more attractive. Consequently, bad agents bargain more cautiously behind closed doors.

[30] A less attractive feature of this model is that if both principals are weak and agents are good, then (not knowing the other agent's type) both would submit a tough bid and hence end up with disagreement. It follows that in this model, disagreement is not the worst outcome for a bad agent. If, on the other hand, the agent's know each other's type, then there would be multiple equilibria.