
A Servant of Two Masters: Communication and the Selection of International Bureaucrats

Leslie Johns

Abstract International bureaucrats must often serve multiple principals who collectively choose policy. How does this affect bureaucrats' incentives to truthfully reveal their private information? I construct a cheap talk model in which a bureaucrat possesses private information about how policies translate into outcomes. The bureaucrat can communicate publicly observable messages about this information to two policymakers, who must then bargain over a set of policy choices. I find that both the bureaucrat's willingness to communicate informatively and the choice of an optimal bureaucrat are highly contingent on the bargaining powers of the two policymakers. When each policymaker is bound to adhere to the bargaining outcome, "moderate" bureaucrats are most preferred. In contrast, when at least one policymaker can leave the bargaining table and exercise an outside option, "biased" bureaucrats can be optimal. I illustrate my findings by examining UN weapons inspections in Iraq from 1991 to 2003.

In November 1997, the United Nations Special Commission (UNSCOM) met to discuss the progress of weapons inspections in Iraq and to deliberate on a forthcoming report to the UN Security Council. It had been a contentious year. Russia, China, and France were pushing for an end to weapons inspections and UN sanctions on Iraq, which had recently expelled all U.S. weapons inspectors. During the meeting, Gennady Gatilov, a Russian member, argued that the draft report was biased, Iraq had disarmed, and UNSCOM should cease its inspection activities. An argument ensued over how the findings of the weapons inspectors should be presented to the Security Council. Three years later, UNSCOM head Richard Butler wrote: "The result was a report as close to the facts as diplomacy would allow."¹

I wish to thank Bruce Bueno de Mesquita, Michael Gilligan, Catherine Hafer, Matias Iaryczower, Dimitri Landa, Adam Meirowitz, Alastair Smith, and the editor and two anonymous reviewers for their feedback. Previous drafts of this article were presented in the New York University Politics Department Research Workshop and the 2006 Annual Conference of the Midwest Political Science Association.

1. Butler 2000, 108.

The political pressure placed on reports by Butler and UNSCOM was not isolated to this one incident. When UNSCOM was created at the end of the Gulf War in 1991, it initially had broad support from the Security Council. However, as Russia, China, and France came to oppose economic sanctions in the late 1990s, they accused UNSCOM of being too biased in favor of U.S. policy. The agency was abolished in 1999 and replaced with the UN Monitoring, Verification, and Inspection Commission (UNMOVIC), which initially had U.S. support. However, after the terrorist attacks of 11 September 2001, the United States grew willing to use military force to disarm Iraq, and U.S. leaders believed that UNMOVIC was too biased against U.S. policy. Throughout the weapons inspection process, changes in the preferences of key UN Security Council members affected both the communication between the Security Council and chief weapons inspectors, and the selection of top-level inspectors.

The situation of Butler is not unique. Governmental and organizational bureaucrats often serve as informational agents: collecting information and reporting back to their superiors about what they have learned. An agent's incentive to report truthfully is affected by how the information that he provides will be used. If the principal will make a decision that the informational agent finds distasteful, then it is unlikely that the agent will tell the truth. In the realm of international politics, many types of individuals can be classified as informational agents, including the heads of many UN agencies, technical advisors at the World Bank, and the mediators and fact-finding commissions that are often used to resolve international conflicts.² This prompts two important substantive questions. First, how do political processes affect the incentives of international bureaucrats to communicate truthfully with their principals? Second, what types of bureaucrats will be most preferred by political actors?

I construct a cheap talk model in which an informational bureaucrat (the "sender" or "agent") possesses private information about how international policies translate into outcomes.³ This bureaucrat has the opportunity to costlessly communicate with two policymakers with heterogeneous preferences (the "receivers" or "principals"). All communication is public—the bureaucrat is unable to speak privately to each policymaker. After the policymakers hear the bureaucrat's message, they bargain over which policy to implement. Once this policy is chosen, the true state of the world is revealed and players receive their final payoffs.

In the traditional cheap talk framework with a single receiver, it is generally the case that both the receiver and sender are better off the more closely their preferences are aligned: increasing the similarity of the players' preferences means that the sender has a greater incentive to send more precise signals, thereby communicating more information in equilibrium. This reinforces the findings of Calvert

2. Kydd 2003.

3. See Crawford and Sobel 1982; and Farrell and Rabin 1996 for an overview of the cheap talk framework; and Osborne 2004, 343–49 for textbook-level introductions to solving cheap talk games.

and Kydd, both of whom argue that it is often rational for a politician to choose advisors who are biased in his favor.⁴ However, it is not clear whether this finding holds with multiple receivers: do international policymakers prefer “biased” informational bureaucrats whose preferences are closely aligned to their own, or do receivers prefer “moderate” bureaucrats whose ideal policies lie midway between the ideal policies of the policymakers?

I find that the answer to this question is highly dependent on certain aspects of the bargaining process. When all parties adhere to the result of the bargaining process (either because they lack the ability to act unilaterally or because it is undesirable for them to do so), then moderate bureaucrats are preferred. That is, each policymaker *A* would rather have an informational agent with moderate preferences than an agent who is biased in *A*'s favor. However, when one player possesses an attractive outside option and is able to leave the bargaining table and unilaterally implement a desired outcome (such as abandoning a UN Security Council resolution and attacking another country in an international interaction), then all policymakers prefer a bureaucratic agent who is biased in favor of the policymaker with the outside option. That is, if country *A* has an attractive outside option, then countries *A* and *B* are best off choosing an international bureaucrat with preferences that are closely aligned with *A*. In both situations the international policymakers want to hire an agent who will provide as much information as possible. However, the identity of this optimal agent depends on features of the international bargaining environment.

Using a cheap talk framework, Milner and Rosendorff show that it is often optimal for a domestic legislature that must ratify an international agreement to value the information provided by domestic interest groups because such information indirectly constrains the set of agreements that an executive can successfully negotiate.⁵ Other international relations scholars have long acknowledged that international organizations can serve as forums in which representatives of nation-states are able to communicate private information about their policy preferences.⁶ Additionally, Morrow conceives of institutions as a way to address the informational problems inherent in the enforcement of international cooperation, and Kydd has highlighted the informational role played by international mediators.⁷ However, little attention has been paid to the role of international bureaucrats in providing information about how international policies translate into outcomes. Additionally, to my knowledge, no systematic theory has been developed regarding the selection of informational bureaucrats by international policymakers. Nonetheless, international organizations, such as the United Nations, World Bank, and International Monetary Fund (IMF), have vast bureaucracies that are charged with

4. See Calvert 1985; and Kydd 2003.

5. Milner and Rosendorff 1996.

6. Keohane 1982.

7. See Morrow 1994; and Kydd 2003.

providing policy-relevant information to member-states, and the individuals who head these bureaucracies often have great influence over policy outcomes.

Previous applications of the cheap talk modeling framework have been limited by their assumption that only one receiver is able to hear the message of the sender.⁸ While invocation of the median voter theorem may make this an innocuous assumption for some contexts, such as a legislative vote, this framework is unsuitable for many political questions, including the study of most international organizations. For example, when an international bureaucrat, such as the head of the International Atomic Energy Agency, reports to the UN Security Council, the message is heard by all members of the Council. The veto power given to permanent members of the UN Security Council ensures that the median voter theorem is an inappropriate framework for modeling such political decision making. To the extent that similarity of preferences leads permanent member-states to form distinct groups, it would be more accurate to describe the policymaking process as a bargaining game in which each faction has the power to veto any particular proposal.⁹

The analysis below is very general in its description of institutional features, such as the bargaining process between policymakers and the availability of outside options. The model has been crafted in this manner to ensure that the results apply across many different institutional contexts and are as general as possible. After characterizing equilibrium behavior, I illustrate my findings by examining the role of UN weapons inspectors in Iraq from 1991 to 2003.

The Model

Suppose there are three players: a sender (S) and two receivers (A and B). The sender privately observes the true state of the world, ω , which is uniformly distributed over the unit interval, $\omega \sim U[0, 1]$.¹⁰ The sender then chooses a message, m , that she communicates to the two receivers. This message conveys information about the true state of the world. I assume that the language in which the sender communicates is sufficiently rich that she is able to convey any meaning that she wishes (the standard “rich language assumption”). For example, S can always find a message to communicate the meaning “ $\omega = 0.3$.” Similarly, S can send messages that the state of the world lies in a particular interval (for example, “ $\omega \in [0.2, 0.4]$ ”) or group of intervals (for example, “ $\omega \in [0.2, 0.4] \cup [0.5, 0.7]$ ”). Math-

8. A rare exception is Farrell and Gibbons 1989, in which a sender communicates with two receivers who then independently choose separate actions. They are not bound to jointly choose an action, as in my framework and most political applications. Similarly, Ainsworth and Sened 1993 construct a model in which a lobbyist signals information to both a unitary actor government and potential members of an interest group. However, final public policy is made by the government alone.

9. The bargaining framework below is similar to the bargaining game used by Voeten 2001 in his study of decision making within the UN Security Council. However, Voeten does not examine how the bargaining process affects communication by informational agents, which is the focus of my analysis.

10. Model notation is also displayed in Table 1.

TABLE 1. Model notation

Variable	Description
ω	True state of the world.
m_j	Message of the sender, indexed by j .
y	Policy collectively chosen by the receivers.
b_i	Bias of player i .
$g(\omega m_j)$	Posterior beliefs of the receiver about the value of the state ω after hearing message m_j .
(x_A, x_B)	Ordered pair denoting the bargaining agreement.
d_i	Disagreement payoff for receiver i .
ϕ_i	The value to receiver i of his outside option.
p_i	The policy that gives receiver i a payoff equivalent to the value of his outside option.
\bar{p}_i	The distance between p_i and receiver i 's induced ideal policy.

ematically, $m \in M = \{X|X \subseteq [0, 1]\}$. Substantively, this rich language assumption restricts attention to the issue of credibility of communication, rather than to issues of meaning.

After hearing the sender's message, the two receivers engage in a bargaining game to choose a collective policy $y \in \mathfrak{Y}$. After the policy is chosen, the true state of the world is observed and payoffs are obtained according to the standard quadratic-loss utility function, $U_i(y|\omega, b_i) = -(y - (\omega + b_i))^2$ where $(\omega + b_i)$ is player i 's ideal point and b_i denotes player i 's bias for $i \in \{S, A, B\}$. I place no restrictions on the value of b_S and assume without loss of generality that $b_A = 0$ and $b_B > 0$. Note that each player's most preferred policy is induced by his beliefs about the true state of the world. Since the sender actually observes ω , her most preferred policy is a function of what she observes about the true state. The receivers cannot observe ω directly but can make inferences about its value based on the sender's message. So each receiver's most preferred policy is induced by beliefs about the value of ω , and these beliefs can be affected by communication with the sender.

The equilibrium concept employed is a perfect Bayesian equilibrium. This requires that each player's actions are sequentially rational given her beliefs, and that these beliefs are consistent with Bayes's rule when possible.¹¹ As Crawford and Sobel show, in equilibrium S will send one of a finite number of messages, each of these messages will indicate a particular interval in which ω lies, and these intervals will form a partition of the state space, $[0,1]$.¹² Such a partition can be described in the following way: if J messages are sent in equilibrium, let

11. Fudenberg and Tirole 2000. The concept of a perfect Bayesian equilibrium places no restrictions on off-the-equilibrium path beliefs. It is trivial to construct these beliefs such that the equilibrium behavior below is supported.

12. Crawford and Sobel 1982.

m_j denote the message sent when $\omega \in [\hat{\omega}_{j-1}, \hat{\omega}_j)$ and $j \in \{1, \dots, J - 1\}$, or $\omega \in [\hat{\omega}_{j-1}, \hat{\omega}_j]$ and $j = J$.¹³ I now examine what the receivers will believe about the value of ω after hearing the sender's message. Since the receivers initially believe that all values of the state variable are equally likely (because ω is uniformly distributed), after hearing that the true state lies in a particular interval, the receivers will believe that all points within that interval are equally likely values. Mathematically, the posterior density function of ω after hearing message m_j is

$$g(\omega|m_j) = \begin{cases} \frac{1}{\hat{\omega}_j - \hat{\omega}_{j-1}} & \text{for } \omega \in [\hat{\omega}_{j-1}, \hat{\omega}_j]; \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Each receiver's expected utility from policy y after hearing message m_j is

$$E[U_i(y|m_j, b_i)] = \int_0^1 U_i(y|\omega, b_i)g(\omega|m_j) d\omega \quad \text{for } i = A, B$$

The sender knows the true state, ω , and in equilibrium she can infer the policy outcome that will be collectively chosen by the receivers based on her message, $\hat{y}(m_j)$. So the sender's utility for sending a message m_j is

$$U_S(m_j|\omega, b_S) = -(\hat{y}(m_j) - (\omega + b_S))^2$$

One can now derive the most preferred policy for receiver i after hearing the sender's message:

Claim 1. Receiver i 's most preferred policy is equal to the midpoint of the message interval plus i 's bias¹⁴:

$$y_i^*(m_j, b_i) = \frac{\hat{\omega}_j + \hat{\omega}_{j-1}}{2} + b_i$$

Recall the assumption that $b_A = 0$ and $b_B > 0$. So player A 's most preferred policy is the midpoint of the interval for which message m_j is sent. In contrast, player B 's most preferred policy is b_B units higher. Now it is possible to examine

13. The precise specification of behavior at the cutpoints (that is, whether the message spaces are open or closed intervals) is irrelevant for deriving equilibrium behavior because the state space is continuous.

14. Proofs for all claims and propositions are provided in the Appendix.




FIGURE 1. *The bargaining region without outside options*

how different bargaining processes affect communication by the sender. I proceed by first considering the results of the model when neither receiver has an outside option in the bargaining game. Then I examine how the possibility of an outside option affects equilibrium behavior.

Cheap Talk and Bargaining Without Outside Options

Bargaining and communication. Figure 1 shows the set of policies that lies in the interval between the most preferred policies of the receivers after hearing a message m_j ,

$$[y_A^*(m_j, b_A), y_B^*(m_j, b_B)] = \left[\frac{\hat{\omega}_j + \hat{\omega}_{j-1}}{2}, \frac{\hat{\omega}_j + \hat{\omega}_{j-1}}{2} + b_B \right]$$

Any policy that is chosen by the two receivers must lie in this interval, and once such a policy is chosen neither receiver can be made better off without making the other receiver worse off. One can mathematically describe any point in this interval by an ordered pair (x_A, x_B) where x_i denotes the extent to which the final negotiated policy choice conforms to player i 's most preferred policy. A higher value of x_A pushes the chosen policy closer to receiver A's most preferred policy, and a higher value of x_B means that the bargaining outcome is closer to B's most preferred policy. This allows one to mathematically define the set of possible bargaining agreements as: $X = \{(x_A, x_B) | 0 \leq x_A \leq b_B \text{ and } x_B = b_B - x_A\}$.¹⁵ The collectively chosen bargaining agreement, $(\hat{x}_A, \hat{x}_B) \in X$, translates into a final policy choice, \hat{y} , in the following way: $\hat{y} = y_A^*(m_j, b_A) + \hat{x}_B$.

15. Muthoo 1999 provides an excellent introductory overview to bargaining games. Exponential elements of this analysis are similar to the framework that he uses.

Since each receiver's utility function is defined over policy outcomes, one must define a new function $V_i(x_i)$ to denote player i 's expected utility from implementing the bargaining agreement characterized by x_i :

$$V_A(x_A) = E[U_A(\hat{y} = y_A^*(m_j, b_A) + b_B - x_A | m_j, b_A)]$$

$$V_B(x_B) = E[U_B(\hat{y} = y_A^*(m_j, b_A) + x_B | m_j, b_B)]$$

If players A and B are unable to agree to a final chosen policy, then each receives a disagreement payoff d_i . Substantively, this payoff could represent player i 's utility from a status quo policy that remains in place if no new policy is implemented.¹⁶ The payoff d_i could also represent utility from situations in which there is no status quo policy, including new policy issues that have not previously been addressed or situations in which a former policy has "expired" because of explicit sunset provisions or an inability/unwillingness to enforce previous policy decisions. I assume that each player has the same disagreement payoff, which is worse than his least-preferred policy from the bargaining region: $d_A = d_B < \min\{V_A(x_A = 0), V_B(x_B = 0)\}$.¹⁷ Substantively, this assumption means that each player would rather accept any agreement from the bargaining region than choose no policy at all.

I assume throughout my analysis that the bargaining game comes from the broad class of protocols that result in the Nash bargaining solution (NBS).¹⁸ The NBS characterizes the outcomes of a diverse set of bargaining protocols, including the infinitely repeated Rubinstein bargaining game when the discount factors of the players converge to one.¹⁹ This is an appropriate solution concept for international negotiations that are controlled by a small number of states with effective veto power, such as deliberations in the UN Security Council, particularly since such negotiations generally lack formal rules that could constitute a binding bargaining

16. Under this interpretation, the players know the utility that they derive from an existing status quo policy, but are uncertain about their utility from implementing a new policy.

17. This technical assumption is needed so that each player has incentive to choose a policy rather than to delay the bargaining process forever. It is analytically equivalent to "forcing" the players to choose a policy. Allowing for $d_A \neq d_B$ affects the specific form of the bargaining outcome but does not affect the substance of my results. Note that the disagreement payoff is not state-contingent: the players know the utility that they derive from a failure to reach agreement.

18. Nash 1950.

19. See Rubinstein 1982; and Binmore 1987. One of the axioms on which the NBS is based is symmetry. Symmetry stipulates that if two players are in identical strategic situations, then the bargaining outcome should give them equivalent payoffs. However, this does not require that players must be in identical strategic situations. It simply means when two players are in identical circumstances, player A does not get a different payoff than player B simply because he is called "player A ." Note that use of the NBS ensures that my results are more robust than if I were to assume a specific extensive form bargaining game.

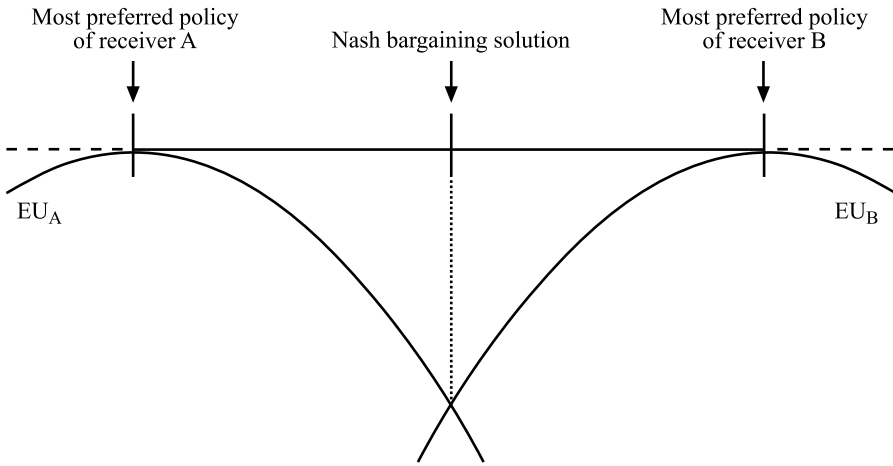


FIGURE 2. Bargaining outcomes without outside options

protocol.²⁰ As shown in Figure 2, the NBS yields the following equilibrium bargaining behavior:

Claim 2. The chosen bargaining agreement after hearing message m_j will be the midpoint of the bargaining interval, and the final chosen policy will be equal to player A's most preferred policy plus half of player B's bias:

$$\hat{y}(m_j) = \frac{\hat{\omega}_j + \hat{\omega}_{j-1}}{2} + \frac{b_B}{2}$$

So the receivers will “split the difference” by choosing a policy that is equidistant between the most preferred policies of the two receivers. Inspection of the partial derivatives of the final chosen policy, $\hat{y}(m_j)$, yields the following comparative statics. As the bias of player B grows, the bargaining region expands and a higher policy is chosen in negotiations. Additionally, when the sender uses higher messages (that is, messages with larger values of $\hat{\omega}_j$ and $\hat{\omega}_{j-1}$), he signals higher

20. Such rules might include restrictions on debate, proposal power, or agenda-setting. This point is also emphasized in Milner and Rosendorff 1996. For scholars interested in institutional contexts for which a bargaining protocol may be well defined and binding, applying alternative bargaining frameworks is a simple modeling extension. The discussion section below examines the robustness of results to such changes.

possible values for the true state of the world, ω . These messages induce the choice of a higher policy.²¹

Now one can consider the sender's incentives for informative communication. An equilibrium becomes "more informative" as more distinct messages are sent by the receiver, which leads to a finer partition of the state space. As with all uses of the cheap talk framework, one can establish an upper bound on how much information can be communicated in equilibrium.²²

Claim 3. Equilibrium communication becomes more informative as the sender's bias approaches half of receiver B's bias (that is, as $b_S \rightarrow (b_B/2)$).

Choosing a sender. Since the willingness of the sender to communicate informatively depends on the bargaining outcome and the sender's bias, what type of sender will be most preferred by the receivers? Cheap talk models with a single receiver generally conclude that optimal informational agents are senders whose preferences are closely aligned with the preferences of the receiver.²³ However, it is not clear that this finding holds when the message sent by the agent induces a bargaining game between multiple principals. I begin by establishing the following counterintuitive finding:

Proposition 1. When receivers cannot exercise outside options, all individuals have identical ex ante preferences over the choice of a sender.

Each receiver's *ex ante* expected utility over possible types of sender (that is, the expected utility of values of the sender's bias, b_S , before learning any information about the true state of the world) can be expressed as a linear transformation of the utility of the other receiver. This means that both utility functions represent the same preferences over the value of the sender's bias, b_S . This result even holds for unmodeled strategic actors who have utility functions of the form $U_i(y|\omega, b_i) = -(y - (\omega + b_i))^2$ over possible policy outcomes but have no meaningful role in the policymaking process. By considering how *ex ante* expected utility for an individual changes in response to changes in the value of the sender's bias, b_S , I establish the following result:

Proposition 2. When receivers cannot exercise outside options, individuals maximize their ex ante expected utility by choosing a "moderate" sender, whose bias is midway between the bias of the two receivers.

21. These are fairly intuitive findings, which suggests that I have selected a bargaining framework that is appropriate for modeling my broader concern: how the presence of a bargaining game affects both communication by an informational agent and the selection of optimal agents.

22. I restrict attention to the most informative equilibrium (that is, the one in which the largest number of distinct and informative messages are sent in equilibrium) because, as shown in the proof of Proposition 2, this equilibrium maximizes the *ex ante* expected utility of all players.

23. For example, Crawford and Sobel 1982.

Both receivers maximize their *ex ante* expected utility by choosing a sender who can provide as much information as is possible. This means that the receivers are best off choosing an agent who will send the largest number of messages possible in equilibrium. As a result of Proposition 1, this holds even for strategic actors who have no influence over the actual bargaining process and so are not explicitly included in the model. In an institution, such as the United Nations, one might believe that certain actors, such as the five permanent members of the Security Council, have great influence over the choice of policies while other actors possess little or no power over policy outcomes. Propositions 1 and 2 imply that these latter actors, who are insignificant in the bargaining process, will have preferences that are identical to those of the major powers in the selection of informational bureaucrats. This means that it does not matter who chooses the sender or how the selection process works because everyone prefers the same type of agent.

If the sender is an extremist, in the sense of having either high or low bias, the sender has an incentive to send vague messages that provide little information. In contrast, a moderate sender, whose most preferred policy is close to the final policy that will be chosen collectively by the receivers, has an incentive to provide very precise information to the committee. This means that each receiver would rather hire a moderate sender than another agent who is biased in the receiver's favor because all policymakers are best off in expectation by learning as much information as is possible.

Cheap Talk and Bargaining with Outside Options

Bargaining and communication. I now consider how the ability of one or more players to “leave the bargaining table” affects communication with an informational agent, as well as the choice of an optimal agent. In bargaining theory, an outside option is the utility that a player can achieve by exiting the bargaining game and engaging in another activity. In a political interaction, one might interpret an outside option as the utility that a player gets from unilaterally imposing a policy. For example, a country may unilaterally impose a tariff, thereby affecting the profits of its firms and of its international trading partners. However, use of an outside option does not necessarily imply that other players are bound to a particular outcome. For example, the U.S. decision to abandon UN weapons inspections in 2003 and proceed with military action did not mean that other UN members were forced to invade as well. In order to make my results as general as possible, I allow both receivers to exercise outside options. Allowing only one player to have an outside option does not affect my findings. Similarly, attaching a specific substantive interpretation to the outside option does not affect my results.

I continue to assume the basic setup from above and let ϕ_i denote the value to receiver i of an outside option that i can exercise. Suppose that $\phi_i < 0$ for $i = A, B$; so each player can conceivably be made better off under a bargaining

agreement than under the outside option. Since outside options are expressed in terms of their utility, one must identify the policy that gives each receiver a value equivalent to his outside option. As shown in Figure 3, I define p_A as the maximum policy that gives player A the same utility as his outside option and \bar{p}_A as the distance between p_A and the most preferred policy of player A that is induced by message m_j . Mathematically, $p_A \equiv \max\{E[U_A^{-1}(\phi_A|m_j, b_A)]\}$ and $\bar{p}_A \equiv p_A - y_A^*(m_j, b_A)$. Similarly, p_B is the minimum policy that gives B the same utility as his outside option, and \bar{p}_B is the distance between p_B and the most preferred policy of player B. So $p_B \equiv \min\{E[U_B^{-1}(\phi_B|m_j, b_B)]\}$ and $\bar{p}_B \equiv y_B^*(m_j, b_B) - p_B$. This means that player A will accept only policies that lie between his most preferred policy and policy p_A , while player B will accept only policies that lie between p_B and his most preferred policy.

I assume throughout that there exists at least one negotiated policy that makes each receiver better off than he would be by unilaterally exercising his outside option. This is equivalent to assuming that $p_B \leq p_A$, which ensures that there exists a bargaining region between policies p_B and p_A . As such, the outside option constricts the bargaining space rather than directly affecting the utility functions of

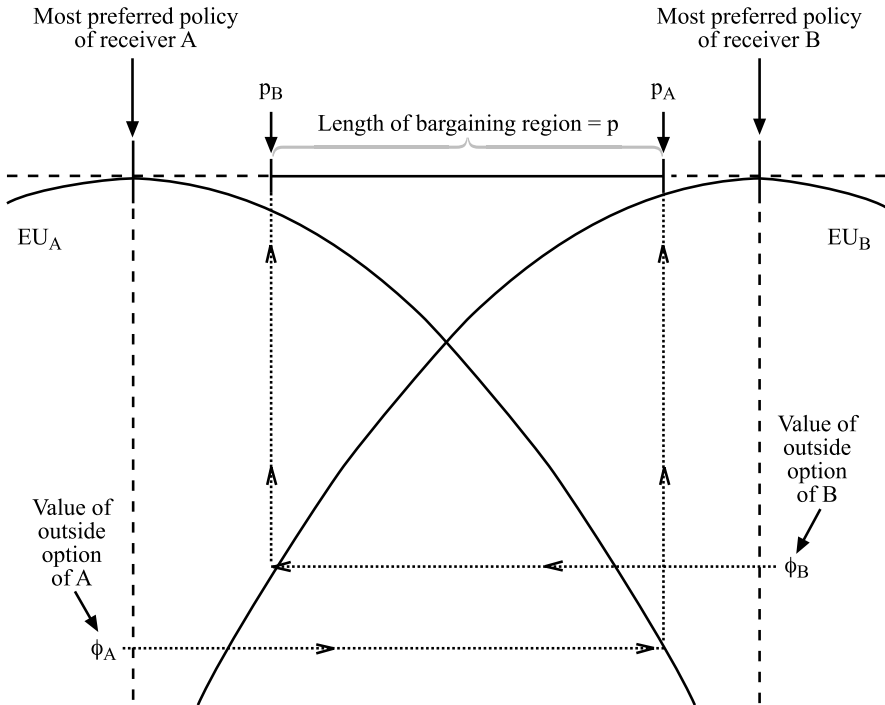


FIGURE 3. Bargaining outcomes with outside options

the players.²⁴ One can continue to describe a bargaining agreement by the ordered pair (x_A, x_B) where x_i denotes the extent to which the chosen policy conforms to player i 's most preferred policy. Let p represent the distance between policies p_A and p_B . So the revised bargaining problem is to select a pair:

$$(\hat{x}_A, \hat{x}_B) \in X = \{(x_A, x_B) | 0 \leq x_A \leq p \text{ and } x_B = p - x_A\}$$

The collectively chosen bargaining agreement translates into a final policy \hat{y} in the following way: $\hat{y} = p_B + \hat{x}_B$.

Since equilibrium bargaining behavior is contingent on the value of each receiver's outside option, I describe receiver A as having a "strong outside option" if receiver A would rather exercise his outside option than accept the "split the difference" policy that would result in a game without outside options.²⁵ One can now proceed to characterize equilibrium bargaining behavior:

Claim 4. If neither receiver has a strong outside option, the final chosen policy will be equal to player A 's most preferred policy plus half of player B 's bias:

$$\hat{y}(m_j) = \frac{\hat{\omega}_j + \hat{\omega}_{j-1}}{2} + \frac{b_B}{2}$$

If receiver A has a strong outside option, the final chosen policy will be the policy that gives player A value equivalent to his outside option, p_A .


If neither receiver has a strong outside option, the bargaining region induced by the outside option, $[p_B, p_A]$, contains the policy that is chosen when there are no outside options. When this occurs, the receivers will "split the difference" by choosing a policy midway between their two most preferred policies. This is shown visually in panel (a) of Figure 4. Substantively, one would expect this to occur when one or both receivers can exercise an outside option, but doing so provides a low level of utility. In contrast, if receiver A 's outside option becomes more desirable than the "split the difference" policy, then the bargaining region will lie completely below the midpoint between the two receivers' most preferred policies. This is displayed in panel (b) of Figure 4. When this occurs, the players no longer split the difference: the chosen policy is the right-most point of the bargaining interval, p_A , which gives receiver A the same amount of utility as he gets from exercising his outside option.²⁶

24. This is the standard formalization of outside options. See Muthoo 1999, 99–105.

25. Note that the assumption that $p_B \leq p_A$ ensures that it is impossible in this model for both receivers A and B to have strong outside options. Since the labeling of receivers is arbitrary, I restrict attention without loss of generality to cases in which A has a strong outside option.

26. This is a manifestation of the outside option principle. See Muthoo 1999, 103.

(a) Neither receiver has a strong outside option



(b) Receiver A has a strong outside option




FIGURE 4. Bargaining outcomes with outside options

Inspection of the partial derivatives of the final chosen policy, $\hat{y}(m_j)$, yields the following comparative statics. The collectively chosen policy, $\hat{y}(m_j)$, is: decreasing in player A's outside option, ϕ_A , when player A has a strong outside option; and increasing in player B's bias, b_B , and the cutpoints associated with the message of the sender, $\hat{\omega}_j$ and $\hat{\omega}_{j-1}$, when neither player has a strong outside option. This means that as the value of the outside option increases for receiver A, he is able to extract an outcome that is closer to his most preferred policy. However, this comparative static does not hold strictly. Increases in player A's outside option only affect the final policy when player A has a strong outside option.

One can now consider how this bargaining behavior affects communication by the sender.

Claim 5. If neither receiver has a strong outside option, equilibrium communication becomes more informative as the sender's bias approaches half of receiver B's bias (that is, as $b_S \rightarrow (b_B/2)$). If receiver A has a strong outside option, equilibrium communication becomes more informative as the sender's bias approaches the distance between player A's ideal policy and the policy equivalent of A's outside option (that is, as $b_S \rightarrow \bar{p}_A$).

A key point to note is that the value of the outside options, ϕ_A and ϕ_B , is constant. This means that the sender's message affects the location of the most preferred policies of the receivers but does not affect the relative positions of the policy equivalents of the outside options (p_A and p_B). So exactly one of two cases must hold across all possible messages that the sender could make: either neither player has a strong outside option and the bargaining region, $[p_B, p_A]$, contains the "split the difference" solution, or receiver A has a strong outside option and the region lies completely below this solution. In the first case, cheap talk behavior is equivalent to sender behavior in the model without outside options. In contrast, when player A has a strong outside option, the bargaining region lies below the midpoint of the receiver's most preferred policies, as in panel (b) of Figure 4. When this occurs, more messages are possible as the sender's bias, b_S , moves toward the distance between player A 's ideal policy and the policy equivalent of A 's outside option, \bar{p}_A .

Choosing a sender. One can now consider the preferences that receivers have regarding the choice of an informational agent. I begin by establishing that the result in Proposition 1 continues to hold, even when one allows for a bargaining game with outside options:

Proposition 3. When receivers can exercise outside options, all individuals have identical ex ante preferences over the choice of a sender.

This means that it is not necessary to explicitly model the process by which the sender is chosen. Since all players have identical *ex ante* preferences over the value of b_S , it does not matter who chooses the sender. One can now consider how different aspects of the bargaining process will affect the choice of an optimal sender:

Proposition 4. When neither receiver has an outside option that is preferred to a policy that "splits the difference" between the two players, then the receivers maximize their ex ante expected utility by choosing a "moderate" sender, whose bias is midway between the bias of the two receivers.

This result follows directly from the combination of Claim 4 and Proposition 2. The analysis above ensures that the presence of outside options does not automatically affect the bargaining process. If neither receiver's outside option is preferred to splitting the difference, then the bargaining and cheap talk behavior is equivalent to equilibrium behavior in a model without outside options. Since the final bargaining outcome will be the midpoint of the two receivers' most preferred policies, a sender who is moderate will be able to communicate more informatively than a sender with low or high bias. Since the political principals are best off when they receive as much information as is possible, each prefers to hire a moderate sender. However, when outside options become sufficiently attractive

that they affect the final bargaining outcome, the optimal choice of a sender will change.

Proposition 5. *If receiver A has a strong outside option, then both receivers maximize their ex ante expected utility by choosing a sender whose preferences are biased toward those of the receiver with the strong outside option.*

If receiver A has an attractive outside option, then the induced bargaining region lies completely below the midpoint of the receivers' most preferred policies. This means that the receivers no longer split the difference: they choose a policy at the right endpoint of the bargaining region induced by the outside options. A sender has incentive to communicate more informatively as her most preferred policy approaches the policy chosen by the receivers. This means that both receivers would like to choose a sender biased in the direction of the player with a strong outside option to maximize the potential for informative communication in equilibrium. Note that this finding implies that there exist circumstances in which a principal finds it optimal to choose an informational agent who is biased against his preferred policy. If a receiver knows that he will be able to extract little at the bargaining table because his competitor has a strong outside option, he will be best off choosing a sender biased in favor of his opponent.

Discussion

Robustness. The substantive results of the model are robust to relaxing many of the assumptions made to simplify the exposition and analysis above. These include quadratic-loss utility functions, uniform prior beliefs about the true state of the world, and the invocation of the NBS. Suppose that any (or all) of the above assumptions are changed. If one can imagine a hypothetical unitary actor R who would find it optimal to unilaterally choose policy $\hat{y}(m_j)$ after hearing message m_j from the sender, then the collectively chosen policy can be rationalized as the result of optimization by an unmodeled strategic player. If the utility function of this hypothetical player satisfies the basic conditions of Crawford and Sobel, then all of the substantive results of this article regarding communication by the sender continue to hold.²⁷ Additionally, if each receiver i 's utility function over the choice of a sender can be expressed as a linear transformation of R 's utility function over

27. Crawford and Sobel 1982, 1437, 1444. These conditions include: $U_R(\cdot)$ is twice continuously differentiable, $(\partial/\partial y)U_R(y^*, \omega, b_i) = 0$ for some $y^* \in \mathfrak{R}$, $(\partial^2/\partial y^2)U_R(y, \omega, b_i) < 0$, $(\partial^2/\partial y \partial \omega)U_R(y^*, \omega, b_i) > 0$, and an arbitrage and monotonicity condition. See Theorem (1) for the general proof strategy.

b_S , then all of the substantive results regarding the selection of an optimal sender continue to hold.²⁸

It is worth noting that I assume throughout my analysis that the value of each receiver's outside option is fixed. I do not allow for state-contingent outside options, in which the utility that a receiver derives from an outside option depends on the true state of the world. If this were true, then the relative position of the induced bargaining region would change across different sender messages. Suppose that one interprets the outside option in the following manner: each of the receivers has some ability to unilaterally impose a policy, but this ability is constrained. So a receiver i cannot unilaterally impose his most preferred policy but can impose a policy that is \bar{p}_i units away. Another possible interpretation is that each receiver can achieve his most preferred policy via an outside option but suffers some disutility from acting unilaterally.²⁹ Both interpretations of the outside option are consistent with the findings of the model as long as this constraint or disutility is constant across all possible states of the world.

Results. The findings above counter assertions often made in the cheap talk and delegation literature that a principal is best off when he chooses an informational agent who has preferences similar to his own. As this analysis shows, principals maximize their *ex ante* utility when they hire a sender who communicates as much information as is possible in equilibrium. The extent to which a sender is willing to provide such information depends on how closely the most preferred policy of the sender coincides with the final chosen policy, not the similarity of preferences of the sender and the receiver. With a unitary actor receiver, this distinction is not important because the receiver can unilaterally implement his most preferred policy. This means that choosing a sender with similar preferences is equivalent to choosing a sender whose most preferred point is close to the chosen policy. However, when a sender's message induces a bargaining game among multiple players, the distinction made above becomes important. It generates situations in which a receiver is best off choosing a sender who is moderate or is even biased in the direction of another player.

When neither receiver has sufficiently attractive outside options, the midpoint of the bargaining interval is chosen as the final policy. This gives many results in this article the "flavor" of results in the previous literature that invoke the median voter theorem in order to justify the assumption of a unitary actor policymaker. However, in this model there does not exist a pivotal policymaker who is located at the median. In many international political situations: (1) the median policymaker is irrelevant because all bargaining power rests with veto players who need not be median voters; or (2) there does not exist a median policymaker because

28. Crawford and Sobel 1982, 1448–49. See Theorems (3) and (4) for the general proof strategy and note that the linear transformation implies identical preferences for all receivers.

29. Voeten 2001 argues persuasively for the latter view in his analysis of decision making within the UN Security Council.

politicians are engaged in bilateral policy negotiations.³⁰ The model above highlights that the key factor affecting the communication and choice of an informational agent is the relationship between his ideal policy and the final policy that will be chosen and implemented by the policymakers. The relationship between the preferences of the sender and the median international policymaker (if one even exists) is irrelevant if the design of an international institution does not allow the median policymaker to deterministically shape the final outcome.

Any resemblance to median voter results disappears once a policymaker in this game has a strong outside option. It is often noted in international relations that no country is bound by the international policymaking process: all countries possess outside options because they can always refuse to comply with any policy that results from multilateral bargaining. As the model above highlights, this does not necessarily affect the policymaking process. If refusing to comply with a multilateral agreement provides low utility, then one should still expect the choice of relatively moderate policies and informational bureaucrats. However, as the desirability of functioning outside of a multilateral agreement increases for a given country, this will result in both the selection of agreements that this country finds more favorable and the choice of international bureaucrats who are more closely aligned with this country's preferences. This accords with the empirical observation that many international organizations, including the World Bank and the IMF, are often perceived as being dominated by bureaucrats favorable to the views of the United States.

The outside options in this model represent the utility that a country obtains from operating outside of the bargaining process. One way to interpret the influence of the outside option is to say that it allows a country to credibly commit itself to a particular bargaining position. The ability to leave the bargaining process allows the leader of country *A* to credibly say, "If the negotiated agreement isn't at least p_A , then there is no way that I will accept this agreement." As such, outside options can be interpreted as reflecting an asymmetric ability to commit to a given policy position.³¹ Such an asymmetry may result from differences between the two international policymakers with regard to subsequent ratification of the negotiated agreement by a domestic legislature.³² Another (not necessarily contradictory) interpretation is that the outside options reflect the distribution of "power" amongst the players. While the precise definition of "power" is open to considerable debate in the international context, Brooks and Wohlforth emphasize the need for a systematic understanding of "how power asymmetries influence bargaining over and within international institutions."³³ To the extent that international "power" can be interpreted as the ability to pursue one's objectives absent a multilateral

30. As Milner and Rosendorff 1996, 152, emphasize, many negotiations that are ostensibly multilateral in nature are de facto led and controlled by two states.

31. Schelling 1960.

32. Milner and Rosendorff 1996.

33. Brooks and Wohlforth 2005, 519.

institution, this analysis provides intuition regarding the impact of asymmetric power on the information-gathering functions of international organizations and provides a modeling framework for further analysis of such questions.

Most surprisingly, Propositions 1 and 3 establish that one should not expect conflict over the choice of such bureaucrats: when bureaucrats must collect and report private information, all countries are better off with the selection of agents who will provide the most information to their political principals.³⁴ Since all principals will have identical preferences over the choice of an informational agent, the process by which such agents are selected should be irrelevant in determining the choice of an agent, and one should not observe conflict between principals during the selection process. However, the model above does not include many other factors that might affect the choice of international bureaucrats. First, countries may genuinely desire diversity and/or representation, irrespective of how this affects final outcomes. For example, representatives of African and underdeveloped countries may have derived inherent utility from the selection of Kofi Annan of Ghana as UN Secretary General, irrespective of the manner in which he affected policy, because he came from their region. Also, international bureaucrats often possess more power than simply collecting information and reporting back to their principals. Some bureaucrats (such as the UN Secretary General) possess additional diverse powers, including budgetary authority and the ability to use their public position to influence the agenda of international organizations. All of these additional factors will surely influence the selection of bureaucrats and could result in conflict between political principals over the choice of an agent. However, as the model above demonstrates, such conflict cannot be induced by the informational functions of bureaucrats. If one observes countries battling over the choice of an international bureaucrat, one must look elsewhere for an explanation.

Evaluating the Model: “Cowboys” and “Bunny-Huggers” in Iraq

Criteria for Evaluating the Model

When evaluating the model, it is important to begin by examining whether the assumptions of the model are consistent with the strategic situations of interest. This includes asking whether an international bureaucrat has preferences over the policies that result from her report; whether the potential for a multilateral agreement exists; and whether the choice of the NBS is appropriate for the institution under consideration. If these key assumptions hold, one can then proceed to examining whether the model’s observable implications accord with reality. The first

34. Note that this contradicts Milner and Rosendorff 1996, which finds that the presence of informed domestic interest groups can raise the utility of the unitary-actor legislature by constraining (and hence lowering the utility) of the executive and his international negotiating partner.

implication of the model is that agents will not fully and truthfully reveal all of their private information to the principals. Second, the model findings imply that changes in the preferences of the political principals can generate dissatisfaction with an agent who was previously optimal. Such dissatisfaction can manifest itself in the replacement of an agent and/or the restructuring of an agency to change the incentives of bureaucrats. If the preferences of one receiver become more extreme, the bargaining region will expand and the midpoint of the interval will move toward the player whose preferences have changed. This means that an agent whose preferences coincide with the previous policy will now begin to seem too biased in favor of the principal whose preferences remained fixed. Alternatively, if the preferences of the policymakers remain fixed but the value of an outside option increases and becomes sufficiently attractive to a given player, then the bargaining outcome will move toward the ideal policy of that player. So a moderate agent who was previously optimal will be less preferred than an agent biased in favor of the player with a strong outside option. Finally, the model predicts that there should be no conflict between political principals over the choice of an optimal informational agent.

International organizations such as the United Nations, the World Trade Organization (WTO)/General Agreement on Tariffs and Trade (GATT), and the World Bank employ thousands of bureaucrats who provide policy-relevant information to member-states. These bureaucrats rarely elicit public attention, largely because there is little controversy in how they are selected and perform their jobs. This accords with the model's finding that one should not expect political principals to engage in conflict over the choice of informational bureaucrats: all policymakers will have identical preferences over such choices and want to hire the agent who will provide the most information in equilibrium. In the case study below, I examine the selection and communication of UN weapons inspectors in Iraq from 1991 to 2003. These individuals were often in the public spotlight both because of the tremendous importance of their work and because of the controversies and contentiousness often generated by their reports to the UN Security Council. As such, this sets a high threshold for evaluating the model because this is a situation in which initial intuition might suggest that the results of the model do not hold. However, to the extent that one can observe strategic behavior consistent with the model, one can be confident in its explanatory power.

Applying the Model to Weapons Inspections in Iraq

At the end of the Gulf War in 1991, the permanent members of the UN Security Council passed Resolution 687, which required Iraq to disarm in the areas of nuclear, chemical, and biological weapons, as well as to destroy all missiles with a range greater than 150 kilometers. As U.S. Secretary of State James A. Baker noted, there was agreement regarding the need for multilateral action: "We were ... determined to maintain substantial economic and political sanctions against Iraq to restrict its aggressive tendencies [W]e needed all our coalition part-

ners to be with us.”³⁵ Clearly, the potential for multilateral agreement within the Security Council existed and there was a need for informational agents to report on Iraqi compliance.

One can loosely think of the application of the fundamentals of the model to this case in the following manner. The unknown state variable, ω , represents the extent to which Iraq was complying with Resolution 687 and the policy parameter, γ , reflects the rewards and/or punishments imposed by the Security Council in response to Iraqi actions. Each permanent member of the Security Council had different preferences regarding how the level of compliance by Iraq should translate into UN policies. These preferences manifest themselves in the bias parameter, b_i . The disagreement payoff, d_i , reflects the utility that the UN Security Council members derived from the absence of any response to the actions of Iraq, including the end of both weapons inspections and sanctions. While policymakers may have differed regarding the relative desirability of this outcome, all believed that some form of multilateral response was preferable to doing nothing. Finally, the value of the outside option, ϕ_i , represents the utility that each of the permanent Security Council members derived from abandoning the multilateral policymaking process. While the UN Charter is clear in establishing the veto power of the five permanent members of the Security Council, the institution lacks explicit and fixed rules regarding the policy formation process.³⁶ When combined with the small number of permanent members that have equal voting weight and often split into opposing factions,³⁷ this makes the NBS an appropriate solution concept.

The UN Security Council allowed the International Atomic Energy Agency (IAEA) to retain its responsibility for nuclear inspections. However, the Security Council created the UN Special Commission (UNSCOM) to collect information about compliance in the other weapons areas. The Security Council pursued a two-pronged strategy when trying to create an optimal informational agent. First, it hired top-level officials who were believed to be supportive of the views of the permanent members. However, unlike the model above in which the sender's bias is assumed to be common knowledge, there is a limit to how much anyone can infer about the preferences and biases of any given individual. So the Security Council also designed the institution in a manner that led to a tough attitude by UNSCOM agents toward Iraqi compliance. The General Assembly and the UN Secretariat, which operates on the principle of geographical representation, had much more lenient attitudes. So the tough-minded Security Council ensured that the head of UNSCOM reported directly to it, ensuring independence from the UN Secretariat. Additionally, all staff and equipment were directly provided and

35. Baker 1995, 441.

36. Even the detailed study of Bailey and Daws on the customs and norms of procedure in the UN Security Council notes that “agreement on the substance of decisions that will be taken, along with the procedure to be followed, is almost always reached in the course of informal consultations prior to a public meeting of the Council.” See Bailey and Daws 1998, 390.

37. Voeten 2001.

paid for by member states in order to remove UNSCOM from interference by the General Assembly, which controls the UN budget.³⁸ When combined with the technical skills needed by weapons inspectors, this ensured that the United States was the dominant provider of personnel and that UNSCOM agents supported the tough attitudes of the Security Council toward compliance.

These attitudes stood in marked contrast to other UN agencies and often generated interagency conflict. Hans Blix, head of the IAEA from 1981 to 1997, wrote: "To UNSCOM, the IAEA inspectors seemed too much like proper civil servants; to the IAEA, some of the UNSCOM inspectors seemed to act Rambo-style . . . UN people who were in Baghdad for various humanitarian functions called UNSCOM staff 'cowboys' and the latter reciprocated by nicknaming the UN staff 'bunny-huggers.'"³⁹ Why did the Security Council not try to inculcate a tough "cowboy" attitude in the IAEA as well? First, the IAEA is an autonomous international organization that is not under the direct control of the Security Council. As such, the Security Council does not have the power to replace the head of the IAEA or to change the structure of the agency. Additionally, the IAEA's broader institutional mandate may have forestalled attempts to indirectly change institutional culture and replace top-level IAEA inspectors because any attempts to alter attitudes toward Iraq would likely have affected the IAEA's activities in other countries. Finally, taking responsibility for nuclear inspections in Iraq away from the IAEA would have generated additional resentment toward UNSCOM, which was already perceived by many outside of the Security Council as "illegitimate."⁴⁰

Throughout the early 1990s the permanent members of the UN Security Council remained committed to tough sanctions against Iraq and the enforcement of Resolution 687. During this period, UNSCOM was a relatively uncontroversial agency. The potential for multilateral agreement in the UN Security Council clearly existed, and the preferences of UNSCOM weapons inspectors coincided with the tough policy being pushed by the permanent five members of the Security Council.

However, the late 1990s saw a gradual change in the preferences of many of these policymakers. In late 1996, Iraq was allowed to resume the sale of oil in exchange for humanitarian supplies under the Oil-for-Food program. As noted by a subsequent independent inquiry on corruption in this program, the vast bulk of Iraqi trade took place with companies based in Russia and France.⁴¹ As U.S. Secretary of State Madeleine Albright recalled, "There was also a general sense of 'sanctions fatigue' and the feeling in some quarters that Baghdad no longer posed a serious threat."⁴² This fatigue and financial pressure took its toll as Russia and France increasingly came to argue for a more lenient approach toward Iraqi com-

38. See Blix 2004, 21–22; and Butler 2000, 40, 63–66.

39. Blix 2004, 22–23.

40. Butler 2000, 206.

41. Volcker 2005.

42. Albright 2003, 275.

pliance. In terms of the model above, this change in preferences lowered their ideal policy, resulting in a wider bargaining interval and a lower midpoint.⁴³

Butler, head of UNSCOM at the time, was acutely aware of this trend. He later described Russia, France, and China as “the three Iraqi sympathizers among the permanent five” of the Security Council and as Iraq’s “powerful allies.”⁴⁴ Such statements clearly reveal that Butler, as with the other weapons inspectors, had well-defined preferences regarding the policies that resulted from his reports—a key assumption of the model above. While UNSCOM was previously a “moderate” agent because the agency’s preferences were not extreme relative to the tough preferences of the Security Council, the changing preferences of Russia, France, and China ensured that UNSCOM came to be perceived as being too aggressive and biased in favor of U.S. policy. As detailed at the beginning of this article, this change in preferences resulted in intense pressure on Butler regarding his reports to the Security Council, and countries attempted to counter the preferences of UNSCOM’s top bureaucrats by forcing new personnel on the agency. Butler recalls that in 1998, “I had to accede to political pressure to accept a Russian and a Frenchman within [the UNSCOM policy] staff. This was designed to inhibit me, and to some extent it worked.”⁴⁵ This accords with the model’s predictions: because Butler had tough preferences that could not be changed, permanent members attempted to moderate his influence by appointing individuals with opposite tendencies, thereby moving the aggregate views of UNSCOM toward the new midpoint of the bargaining interval.⁴⁶

In January 1999, it was revealed in the Western press that undercover U.S. intelligence officials had been working as UNSCOM weapons inspectors.⁴⁷ This revelation confirmed the suspicions of many that UNSCOM was too biased and sealed UNSCOM’s reputation as an agency full of overly aggressive “cowboys.” By the end of the year, the Security Council passed Resolution 1284, which replaced UNSCOM with a new informational agency: the United Nations Monitoring, Verification and Inspection Commission (UNMOVIC). Because the permanent members of the Security Council now had diverse preferences that more closely mirrored those of the Secretariat and General Assembly, institutional controls were placed on UNMOVIC to ensure that it reflected this change. The executive chairman of UNMOVIC was required to report to the UN Secretary General, removing the previous institutional autonomy of UNSCOM. Also, the new weapons inspection agency was funded by the Oil-for-Food program, rather than by the voluntary

43. The changing preferences of Russia and France regarding the severity of Iraqi sanctions did not affect the value of exercising an outside option because overriding existing UN sanctions cannot be done unilaterally.

44. Butler 2000, 91, 92.

45. *Ibid.*, 132.

46. Note that the model also predicts that the United States should have been supportive of these staff changes. Unfortunately, there is not enough information in the historical record to evaluate whether this finding of the model is supported or contradicted.

47. See *Boston Globe*, 6 January 1999, A1; and *New York Times*, 7 January 1999, A1, A6.

provision of personnel and equipment by nation-states. These measures ensured that UNMOVIC was subject to Article 101 of the UN Charter, which mandates the recruitment of international civil servants on the basis of geographic representation. This ensured that agency was not dominated by U.S. citizens with tough attitudes toward Iraqi compliance. The principals knew that they were best off choosing an agency head who had more lenient preferences than those embodied by UNSCOM. They chose the chief “bunny-hugger”: Blix, the former head of the IAEA. During its first few years, UNMOVIC operated with little controversy and the permanent members of the Security Council appeared to be pleased with its work.

However, the terrorist attacks of 11 September 2001 drastically changed the political landscape. While the United States had never explicitly ruled out unilateral preemptive action against countries perceived as threats, the terrorist attacks of 2001 increased the utility of exercising such an outside option.⁴⁸ The presence of such a strong outside option is a key element of many of the model findings above. The ability of a policymaker to profitably leave multilateral negotiations ensures that the bargaining outcome will move closer to his ideal policy. This change in the bargaining process in turn affects the sender’s incentives for informative communication. The increase in the value of the U.S. outside option after 2001 meant that the anticipated policy chosen in response to Blix’s messages moved farther away from his ideal policy, thereby decreasing his incentive to communicate informatively. Accordingly, as preparations were made in early 2003 for unilateral action against Iraq, Blix’s reports became increasingly cryptic. In late February, the *New York Times* reported: “Asked if the report he delivered to the United Nations on Friday represented a glass half full or half empty, Hans Blix . . . replied, ‘It is a glass in which they have poured more water.’ Mr. Blix’s practice of making Delphic comments like that one and balancing the positive and negative on Iraq has frustrated the White House and strained his relations with Washington.”⁴⁹

It is rare for high-level bureaucrats to admit to making imprecise statements or withholding information from their political principals. Nonetheless, Blix subsequently admitted to doing so: “I said [to the Security Council] that there were ‘strong indications’ that Iraq had produced more anthrax than it declared and that some of it might still exist. I did not go further than that. An expert briefing I had had on this matter had strongly suggested that a quantity of anthrax had been retained when anthrax was being destroyed in 1991, and could still be effective if Iraqis had succeeded in drying the agent. However, the evidence, although impressive, had not been compelling.”⁵⁰ Blix’s admission provides evidence in support of the model’s finding that agents will not fully and truthfully reveal all of their information to their principals.

48. U.S. National Security Council 2002.

49. *New York Times*, 2 March, 2003, 13.

50. Blix 2004, 139–40.

Blix's reports during this period often induced intense conflict amongst his political principals. If the desirability of exercising an outside option by invading Iraq had increased for the United States and Great Britain following the 11 September 2001 terrorist attacks, why did the UN Security Council members not agree to replace Blix with a weapons inspector who was more closely aligned with the preferences of the United States and Great Britain? If the findings of the model hold, such a replacement would have communicated more information to the Security Council in equilibrium and all members of the Security Council should have identical preferences regarding the value of such a replacement. However, it is not likely that such a personnel change would have forestalled military action. Simply stated, the expected payoff from invading Iraq was most likely perceived by U.S. and British policymakers as far superior to any agreement that could be negotiated within the confines of the Security Council. It was no longer true that there existed a multilateral agreement that made all receivers better off than exercising their outside option. One can think about the delay between September 2001 and the abandonment of the weapons inspection process in early 2003 as reflecting factors not present in the model. For example, perhaps Bush and Blair were initially uncertain about the cost of military action for reasons unrelated to the weapons inspection process, thereby explaining their willingness to allow Blix to retain his position until they had considered the feasibility of unilateralism. Also, the need to persuade the British and U.S. public of the desirability of invasion and to build up military forces contributed to the delay. Additionally, perhaps the other UN Security Council members were uncertain regarding the value that the United States and Great Britain derived from abandoning the UN weapons inspection process and acting unilaterally. While none of these factors are explicitly modeled, all are consistent with the basic framework and findings above.

The conflict generated by Blix's reports is not consistent with the model's finding that political principals should have identical preferences over the choice of an informational agent. It is important to note that in the model above there exist two possible ways in which conflict between political principals can be induced: (1) conflict over the choice of a bureaucratic agent; and (2) conflict over the choice of a policy in response to the messages of an agent. While the model analytically demonstrates that the former will not take place in equilibrium, there is still potential for conflict in the unmodeled policymaking process. However, to the extent that conflict in the Security Council resulted from the choice of Butler as an agent, and not simply from attempts by Security Council members to shape policies in response to his reports, one aspect of the model is not supported.

Weapons inspectors, such as Butler and Blix, no doubt had good intentions during the years that they served as international civil servants. However, they were also individuals with preferences regarding the policies that would result from the information that they provided. As high-level bureaucrats they had a great deal of discretion in their communication with their principals. As the model above explains, these two factors affected the both the credibility and the informativeness of their communication. In accordance with the model's findings, changes in

the political preferences of key members of the Security Council affected the choice of informational agents. When Russia and France shifted to a more lenient approach to Iraqi compliance with UN Resolution 687, UNSCOM could no longer serve effectively: a change of informational agent was needed. UNMOVIC, which was designed by all of the permanent members of the Security Council to be more moderate than its predecessor, was able to function effectively for a few years. However, the terrorist attacks of 2001 increased the willingness of the United States to exercise an outside option via unilateral military action. The analysis above explicitly modeled how the presence of such an outside option affects both bargaining between policymakers and communication by an informed agent. In accordance with the model's findings, UNMOVIC was now too moderate to be able to communicate effectively with its political principals.

While this case study provides strong overall support for the model, the conflict induced by Butler's reports in the post-2001 period does not accord with the model finding that all principals will have identical preferences over the choice of an informational agent. There are two ways in which the model might be profitably expanded to account for such conflict. First, the sender's bias, b_S , was assumed to be common knowledge. To the extent that this parameter is private information, conflict over the choice of an optimal agent can be induced in the game. All receivers will agree on what the optimal level of bias is in an agent but may disagree on which potential agent has preferences that match this level of bias. Second, the value of the outside option was fixed and common knowledge in the model above. However, principals may have incentive to argue about the choice of the agent to build support from domestic political audiences. Such support could influence the bargaining process to the extent that it shapes the value of the outside option. For example, to the extent that the U.S. public believed Blix was an Iraqi appeaser, leaving the Security Council and taking unilateral action became more desirable because multilateral action appeared to be ineffective. By explicitly modeling domestic responses to international behavior, the model may be profitably expanded to allow for endogenous outside options.

Conclusions

The cheap talk literature has provided great insight into situations in which an informational agent can communicate with a single receiver. However, these results have limited applicability because informational agents, such as bureaucratic agencies, policy advisors, and international fact-finding commissions, must often serve multiple principals with heterogeneous preferences. As I demonstrate, the presence of multiple audiences affects the incentives for a sender to informatively communicate with his receivers.

I examine strategic communication when a report made by a single informational bureaucrat is heard by two policymakers who then bargain over a set of policy choices. I find that both the bureaucrat's willingness to communicate informatively and the choice of an optimal bureaucrat are contingent on the bargaining

powers of the two policymakers. These international policymakers are best off choosing agents who will provide the most information in equilibrium. A bureaucrat's incentive to communicate informatively increases as her most preferred policy approaches the policy chosen by the policymakers. This differs from previous research, in which the similarity of the preferences of the sender and receiver is argued to be key to informative communication. When each policymaker is bound to adhere to the bargaining outcome, the receivers will choose policies located at the midpoint of their two most preferred policies. Since moderate policies will be chosen, moderate bureaucrats are most preferred *ex ante* by both policymakers. In contrast, when at least one principal can leave the bargaining table and exercise an outside option, incentives can change. If one policymaker has a sufficiently attractive outside option, he is able to extract more in the bargaining game. This means that both principals will prefer to choose a bureaucrat who is biased in the direction of the principal with the attractive outside option.

This theoretical framework is particularly important for the study of international institutions—a research area in which results based on a median voter framework can rarely be invoked. My analysis complements past research on policy-making within the UN Security Council.⁵¹ It also demonstrates that while the findings of Kydd regarding the value of biased mediators in international negotiation may sometimes extend to the choice of international policy advisors,⁵² there are also circumstances in which neutral bureaucrats are optimal. This constitutes a substantive contribution toward understanding the role of bureaucrats within international institutions.

Appendix

Proof of Claim 1.

$$\begin{aligned} \frac{\partial}{\partial y} E[U_i(y|m_j, b_i)] &= \frac{\partial}{\partial y} \int_{\hat{\omega}_{j-1}}^{\hat{\omega}_j} [-(y - (\omega + b_i))^2] \left[\frac{1}{\hat{\omega}_j - \hat{\omega}_{j-1}} \right] d\omega \\ &= -2y + 2b_i + \hat{\omega}_j + \hat{\omega}_{j-1} = 0 \\ \Leftrightarrow y_i^*(m_j, b_i) &= \frac{\hat{\omega}_j + \hat{\omega}_{j-1}}{2} + b_i \end{aligned}$$

Proof of Claim 2. By definition, the NBS is the agreement that solves the following optimization problem:

$$\max_{(x_A, x_B) \in X} (V_A(x_A) - d_A)(V_B(x_B) - d_B).$$

51. Voeten 2001.

52. Kydd 2003.

Note that since the induced utility functions are defined over the set X , $V_i(x_i)$ is strictly increasing in x_i for $i = A, B$. This means that the inverse function $V_i^{-1}(v_i)$ exists over X . So for a given utility level v_i , there exists exactly one bargaining agreement x_i such that $V_i(x_i) = v_i$. When combined with the differentiability and concavity of V_i for $i = A, B$, this ensures that the NBS is unique and solves the following equation:

$$\frac{V_A(x_A) - d_A}{V'_A(x_A)} = \frac{V_B(x_B) - d_B}{V'_B(x_B)} \tag{1}$$

Filling the relevant values into equation (1) yields the following NBS: $\hat{x}_B = (b_B/2)$.

Proof of Claim 3. For a J -message monotonic equilibrium to exist, it must be that case that

$$U_S(m_j|\hat{\omega}_j, b_S) = U_S(m_{j+1}|\hat{\omega}_j, b_S) \quad \text{for } j = 1, \dots, J - 1$$

This implies that for $j = 1, \dots, J - 1$

$$\begin{aligned} & -\left(\frac{\hat{\omega}_{j-1} + \hat{\omega}_j + b_B}{2}\right)^2 + (\hat{\omega}_{j-1} + \hat{\omega}_j + b_B)(\hat{\omega}_j + b_S) \\ & = -\left(\frac{\hat{\omega}_j + \hat{\omega}_{j+1} + b_B}{2}\right)^2 + (\hat{\omega}_j + \hat{\omega}_{j+1} + b_B)(\hat{\omega}_j + b_S) \end{aligned}$$

Given the monotonicity of $\hat{\omega}_j$, this can only hold if

$$\hat{\omega}_{j+1} = 2\hat{\omega}_j - \hat{\omega}_{j-1} + 4b_S - 2b_B$$

Since $\hat{\omega}_0 = 0$, the class of solutions satisfying the difference equation is

$$\hat{\omega}_j = j\hat{\omega}_1 + j(j-1)(2b_S - b_B) \quad \text{for } j = 1, \dots, J$$

Let $N(b_S, b_B)$ denote the largest number of messages that are possible in equilibrium. Then $N(b_S, b_B)$ is the largest positive integer j such that $j(j-1)|2b_S - b_B| < 1$ and $N(\cdot)$ increases as $b_S \rightarrow (b_B/2)$.

Proof of Proposition 1. Let $N = N(b_S, b_B)$. The *ex ante* expected utility of player i from a sender with bias b_S is

$$\begin{aligned} E[U_i(b_S|\cdot)] &= \sum_{j=1}^N \Pr(m_j(b_S))E[U_i(\hat{y}(m_j)|\cdot)] \\ &= \sum_{j=1}^N (\hat{\omega}_j - \hat{\omega}_{j-1}) \int_{\hat{\omega}_{j-1}}^{\hat{\omega}_j} \left[-\left(\frac{\hat{\omega}_{j-1} + \hat{\omega}_j}{2} + \frac{b_B}{2} - (\omega + b_i)\right)^2 \right] \left[\frac{1}{\hat{\omega}_j - \hat{\omega}_{j-1}} \right] d\omega \\ &= b_B b_i - b_i^2 - \frac{b_B^2}{4} - \frac{1}{12} \sum_{j=1}^N \left[\frac{1}{N} + (2j - N - 1)(2b_S - b_B) \right]^3 \end{aligned}$$

Note that $E[U_j(b_S|\cdot)] = \alpha + E[U_i(b_S|\cdot)]$, where $\alpha = b_B(b_j - b_i) + (b_i^2 - b_j^2)$. Since $E[U_j(b_S|\cdot)]$ is a positive linear transformation of $E[U_i(b_S|\cdot)]$, both expected utility functions represent the same preferences over b_S .

Proof of Proposition 2. Let $x = 2b_S - b_B$ and $N = N(b_S, b_B)$. From the proof of Proposition 1,

$$\frac{\partial}{\partial x} E[U_i(b_S|\cdot)] = -\frac{1}{4} \sum_{j=1}^N \left[\frac{1}{N} + (2j - N - 1)x \right]^2 \left[-\frac{1}{N^2} \frac{\partial N}{\partial x} + (2j - N - 1) - x \frac{\partial N}{\partial x} \right]$$

Note that N is the maximum integer j such that $j(j - 1)|x| < 1$. So N is a discontinuous step function. For most values of x , $(\partial N/\partial x) = 0$. For such values of x :

$$\frac{\partial}{\partial x} E[U_i(b_S|\cdot)] = -\frac{1}{4} \sum_{j=1}^N \left[\frac{1}{N} + (2j - N - 1)x \right]^2 (2j - N - 1) = -\frac{1}{6} [(N^2 - 1)x]$$

Note that when $x < 0$, $(\partial/\partial x)E[U_i(b_S|\cdot)] > 0$, and when $x > 0$, $(\partial/\partial x)E[U_i(b_S|\cdot)] < 0$. So as $x \rightarrow 0$, the utility of player i increases. One must now consider what happens at the points of discontinuity of function N . Consider values $x > 0$ and $\hat{N} \geq 2$ such that $\hat{N}(\hat{N} - 1)|x - \varepsilon| < 1$ and $\hat{N}(\hat{N} - 1)|x| = 1$. So an \hat{N} -message equilibrium is possible for $x - \varepsilon$, but not for x . Compare the *ex ante* utility for player i of an \hat{N} -message equilibrium to an $(\hat{N} - 1)$ -message equilibrium, holding x constant:

$$\begin{aligned} & E[U_i(\hat{N}|x, \cdot)] - E[U_i(\hat{N} - 1|x, \cdot)] \\ &= -\frac{1}{12} \left[\sum_{j=1}^{\hat{N}} \left(\frac{1}{\hat{N}} + (2j - \hat{N} - 1)x \right)^3 - \sum_{j=1}^{\hat{N}-1} \left(\frac{1}{\hat{N}-1} + (2j - (\hat{N} - 1) - 1)x \right)^3 \right] \\ &= -\frac{1}{12} \left[\frac{(2\hat{N} - 1)(\hat{N}^2 x^2 (\hat{N} - 1)^2 - 1)}{(\hat{N} - 1)^2 \hat{N}^2} \right] < 0 \end{aligned}$$

So an infinitely small decrease in the value of x raises utility by increasing the number of messages that are possible in equilibrium. Analogous reasoning shows that for values $x < 0$ and $\hat{N} \geq 2$, such that $\hat{N}(\hat{N} - 1)|x| = 1$ and $\hat{N}(\hat{N} - 1)|x + \varepsilon| < 1$, an infinitely small increase in x raises utility by increasing the number of messages that are possible in equilibrium. Furthermore, this shows that all players' *ex ante* utility is maximized in the equilibrium with the finest partition of messages that is possible. This means that player i is best off choosing a sender such that x is as close to 0 as is possible. Also, player i is indifferent between $x' > 0$ and $x'' < 0$ if and only if $-x'' = x'$.

Proof of Claim 4. The derivation of equation (1) in the proof of Claim 2 continues to hold, which yields the following NBS:

$$\hat{x}_B = \begin{cases} p & \text{if } p_A < \frac{\hat{\omega}_{j-1} + \hat{\omega}_j + b_B}{2}; \text{ and} \\ \frac{\hat{\omega}_{j-1} + \hat{\omega}_j + b_B}{2} - p_B & \text{if } p_B \leq \frac{\hat{\omega}_{j-1} + \hat{\omega}_j + b_B}{2} \leq p_A \end{cases}$$

Proof of Claim 5.

Case 1: Suppose $p_A < (\hat{\omega}_{j-1} + \hat{\omega}_j + b_B)/2$. For a J -message monotonic equilibrium to exist, the following must hold for $j = 1, \dots, J - 1$:

$$\begin{aligned} & -\left(\frac{\hat{\omega}_{j-1} + \hat{\omega}_j}{2} + \bar{p}_A\right)^2 + (\hat{\omega}_{j-1} + \hat{\omega}_j + 2\bar{p}_A)(\hat{\omega}_j + b_S) \\ & = -\left(\frac{\hat{\omega}_j + \hat{\omega}_{j+1}}{2} + \bar{p}_A\right)^2 + (\hat{\omega}_j + \hat{\omega}_{j+1} + 2\bar{p}_A)(\hat{\omega}_j + b_S) \end{aligned}$$

Given the monotonicity of $\hat{\omega}_j$, this can only hold if

$$\hat{\omega}_{j+1} = 2\hat{\omega}_j - \hat{\omega}_{j-1} + 4b_S - 4\bar{p}_A$$

Since $\hat{\omega}_0 = 0$, the class of solutions satisfying the difference equation is

$$\hat{\omega}_j = j\hat{\omega}_1 + j(j-1)(2b_S - 2\bar{p}_A) \quad \text{for } j = 1, \dots, J$$

Let $N(b_S, \phi_A)$ denote the largest number of messages that are possible in equilibrium. Then $N(b_S, \phi_A)$ is the largest positive integer j such that $j(j-1)|2b_S - 2\bar{p}_A| < 1$ and $N(\cdot)$ increases as $b_S \rightarrow \bar{p}_A$.

Case 2: Suppose $p_B \leq (\hat{\omega}_{j-1} + \hat{\omega}_j + b_B)/2 \leq p_A$. Then the bargaining game induces sender behavior that is equivalent to cheap talk in the game without outside options.

Proof of Proposition 3. Suppose $p_B \leq (\hat{\omega}_{j-1} + \hat{\omega}_j + b_B)/2 \leq p_A$. Then $\hat{y} = (\hat{\omega}_{j-1} + \hat{\omega}_j)/2 + (b_B/2)$ and the proof of Proposition 1 holds.

Suppose $p_A < (\hat{\omega}_{j-1} + \hat{\omega}_j + b_B)/2$. Then $\hat{y} = (\hat{\omega}_{j-1} + \hat{\omega}_j)/2 + \bar{p}_A$. Let N denote the largest number of messages that are possible in the corresponding cheap talk equilibrium. The following holds:

$$\begin{aligned} E[U_i(b_S|\cdot)] &= \sum_{j=1}^N (\hat{\omega}_j - \hat{\omega}_{j-1}) \int_{\hat{\omega}_{j-1}}^{\hat{\omega}_j} \left[-\left(\frac{\hat{\omega}_{j-1} + \hat{\omega}_j}{2} + \bar{p}_A - (\omega + b_i)\right)^2 \right] \left[\frac{1}{\hat{\omega}_j - \hat{\omega}_{j-1}} \right] d\omega \\ &= 2\bar{p}_A b_i - b_i^2 - \bar{p}_A^2 - \frac{1}{12} \sum_{j=1}^N \left[\frac{1}{N} + (2j - N - 1)(2b_S - 2\bar{p}_A) \right]^3 \end{aligned}$$

Note that $E[U_j(b_S|\cdot)] = \alpha + E[U_i(b_S|\cdot)]$, where $\alpha = 2\bar{p}_A(b_j - b_i) + (b_i^2 - b_j^2)$. Since $E[U_j(b_S|\cdot)]$ is a positive linear transformation of $E[U_i(b_S|\cdot)]$, both expected utility functions represent the same preferences over b_S .

Proof of Proposition 4. If $p_B \leq (\hat{\omega}_{j-1} + \hat{\omega}_j + b_B)/2 \leq p_A$, then $\hat{y} = (\hat{\omega}_{j-1} + \hat{\omega}_j + b_B)/2$, as in the bargaining model without outside options. This ensures results equivalent to Proposition 2.

Proof of Proposition 5. Let N denote the largest number of messages that are possible in the corresponding cheap talk equilibrium and $x = 2b_S - 2\bar{p}_A$. Then the results derived in the proof of Proposition 2 regarding $(\partial/\partial x)E[U_i(b_S|\cdot)]$ continue to hold. As $x \rightarrow 0$, the *ex ante* expected utility of player i increases.

References

- Ainsworth, Scott, and Itai Sened. 1993. The Role of Lobbyists: Entrepreneurs with Two Audiences. *American Journal of Political Science* 37 (3):834–66.
- Albright, Madeleine. 2003. *Madam Secretary*. New York: Miramax Books.
- Bailey, Sydney D., and Sam Daws. 1998. *The Procedure of the UN Security Council*. 3d ed. Oxford, England: Clarendon Press.
- Baker, James A., III 1995. *The Politics of Diplomacy: Revolution, War, and Peace, 1989–1992*. New York: Putnam's Sons.
- Binmore, Ken. 1987. Nash Bargaining Theory II. In *The Economics of Bargaining*, edited by Ken Binmore and Partha Dasgupta, 27–46. Oxford, England: Basil Blackwell.
- Blix, Hans. 2004. *Disarming Iraq*. New York: Pantheon Books.
- Brooks, Stephen G., and William C. Wohlforth. 2005. International Relations Theory and the Case Against Unilateralism. *Perspectives on Politics* 3 (3):509–24.
- Butler, Richard. 2000. *The Greatest Threat: Iraq, Weapons of Mass Destruction, and the Growing Crisis of Global Security*. New York: Public Affairs.
- Calvert, Randall L. 1985. The Value of Biased Information: A Rational Choice Model of Political Advice. *Journal of Politics* 47 (2):530–55.
- Crawford, Vincent P., and Joel Sobel. 1982. Strategic Information Transmission. *Econometrica* 50 (6):1431–51.
- Farrell, Joseph, and Robert Gibbons. 1989. Cheap Talk with Two Audiences. *American Economic Review* 79 (5):1214–23.
- Farrell, Joseph, and Matthew Rabin. 1996. Cheap Talk. *Journal of Economic Perspectives* 10 (3):103–18.
- Fudenberg, Drew, and Jean Tirole. 2000. *Game Theory*. Cambridge, Mass.: MIT Press.
- Keohane, Robert O. 1982. The Demand for International Regimes. *International Organization* 36 (2):325–55.
- Kydd, Andrew. 2003. Which Side Are You on? Bias, Credibility, and Mediation. *American Journal of Political Science* 47 (4):597–611.
- Milner, Helen V., and B. Peter Rosendorff. 1996. Trade Negotiation, Information and Domestic Politics: The Role of Domestic Groups. *Economics and Politics* 8 (2):145–89.
- Morrow, James D. 1994. Modeling the Forms of International Cooperation: Distribution Versus Information. *International Organization* 48 (3):387–423.
- Muthoo, Abhinav. 1999. *Bargaining Theory with Applications*. Cambridge: Cambridge University Press.
- Nash, John F., Jr. 1950. The Bargaining Problem. *Econometrica* 18 (2):155–62.
- Osborne, Martin J. 2004. *An Introduction to Game Theory*. New York: Oxford University Press.
- Rubinstein, Ariel. 1982. Perfect Equilibrium in a Bargaining Model. *Econometrica* 50 (1):97–110.
- Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge, Mass.: Harvard University Press.
- U.S. National Security Council (NSC). 2002. *National Security Strategy of the United States of America*. Washington, D.C.: Government Printing Office.
- Voeten, Erik. 2001. Outside Options and the Logic of Security Council Action. *American Political Science Review* 95 (4):845–58.
- Volcker, Paul A. 2005. Manipulation of the Oil-for-Food Programme by the Iraqi Regime. Report of the Independent Inquiry Committee into the United Nations Oil-for-Food Programme, October 27. Available at www.iic-offp.org/documents/Final%20Report%2027Oct05/IC%20Final%20Report%20-%20Cover%20and%20Index.pdf. Accessed 20 December 2006.