A Formal Theory of the Initiation and Implementation of Mediation in Wars*

Bernd Beber†

June 2008

Abstract

Recent years have seen renewed debate about the merits of mediator characteristics such as impartiality in settling wars. Less attention has been paid to the issue of how mediators become involved in conflicts in the first place. This is an important question to ask, because inferences about the determinants of mediation success can be biased if certain types of mediation are systematically selected out of the observed sample of conflict management events. This paper assesses the issue of strategic mediation initiation in a game-theoretic framework, and the key advantage the paper offers over most existing approaches is that it presents a model and equilibrium analysis of both mediation initiation and implementation.

1 Introduction

Mediation is by some measures the most common type of intervention in international disputes (Bercovitch and Diehl, 1997: 3), and according to some, also one of the most effective (Rauchhaus, 2006; Dixon, 1996: 671).[1] Mediators typically use a variety of non-coercive tactics, but are generally understood to not promise rewards,

*I am indebted to Macartan Humphreys, Adriana Lins de Albequerque, Massimo Morelli, Kris Ramsay, Alex Scacco, Michael Ting, Matt Winters, participants of the Columbia University International Politics Seminar, and panel participants at the 2007 Annual Meetings of the Midwest Political Science Association and the American Political Science Association for valuable comments. Any remaining errors are entirely mine.

†Columbia University, bhb2102@columbia.edu

[1] Mediation is defined as a non-coercive intervention by a third party in a dispute between two or more political entities. See for instance Princen (1992: 3) and Bercovitch (1996: 3). The intervening third party does not have to be a state. It can be, and often is, an international governmental or non-governmental organization, or even an ad hoc group of individuals. The political entities engaged in conflict do not have to be states either, and mediation is a common means of intervention in intrastate violence. I use the terms dispute and war interchangeably, but in either case I refer to high-intensity, militarized conflict.

threaten punishment, or deliver services or goods to the disputants (Skjelsbæk and Fermann, 1996: 76).[2] One explanation for how mediators help settle conflicts in the absence of coercion, and the one this paper will focus on, is that mediators can bring information to the table that the disputants were not aware of. If conflict onset or escalation is the result of uncertainty and disagreement about the gains each disputant can expect to reap, then mediators can help resolve disputes by providing information that alleviates this uncertainty.[3] This is obviously not the only type of mediation we observe: Mediators can manage public relations and shield leaders from domestic audiences, provide procedural advice, meeting places, deadlines, and many other things. But mediation that deals in private information is particularly intriguing because private information has gained currency as a possible explanation for war (Fearon, 1995), and because this type of intervention makes for a particularly cost-effective policy option. NGOs, for example, do not typically have the capacity to change the material capabilities that underpin disputants' beliefs about their strength, but they can gather and disseminate information not previously known to the adversaries. Information-based mediation could be our best shot at combining resource efficiency with effective conflict resolution.[4]

This raises two questions. First, what kind of mediators are most effective at conveying dispute-relevant information and guiding disputants towards a settlement? Second, are those kinds of mediators also the types of mediators that tend to become involved in disputes? This paper provides hypotheses in response to both questions, by way of a game-theoretic analysis of mediation initiation and implementation. Most models do not theorize about both steps in this two-stage process within the same game-theoretic framework.[5] This is important, as some recent statistical work has argued (Beardsley and Schmidt, 2006), because selection mechanisms affect the

[2]Fisher (1995) distinguishes between "pure mediation" in which the mediator has no leverage to use "promised rewards or threatened punishments to motivate the parties toward a settlement" and "power mediation" (or "mediation with muscle"), which is really an instance of "triadic bargaining." This paper focuses on the former type. On different means of conflict management, see also Dixon (1996: 658).

[3]For an early example of an information-based approach to mediation, see Young (1972: 56–7).

[4]As a piece of anecdotal evidence that policy-makers acknowledge the importance of uncertainty in decision-making, consider the account by Gary Sick, National Security Council staff member for Iran under President Carter, of Algeria's mediation during the Iranian hostage crisis, when negotiations were beset by "uncertainty about the private strategies and motivations of third parties or adversaries, and the natural tendency to interpret events from one's own perspective, . . . the same limitations that apply to any participant" (Sick, 1985: 22).

[5]Terris and Maoz (2005) is an exception. Recent models of the implementation of mediation include Kydd (2003, 2006); Smith and Stam (2003); Rauchhaus (2006).

sample of observed mediation events and can distort results we draw from an analysis of that sample.

I also address the issue of equilibrium selection, and provide a complete mapping of all equilibria of a particular class in the game. Previous models have tended to focus only on the conditions under which peaceful equilibria can be sustained, which does not directly address whether there are regions of the parameter space in which either peace or continued conflict are unique equilibrium outcomes, and if there are multiple equilibria, why one might prevail. Furthermore the paper allows for a wide range of mediator types and is not limited to an analysis of a small set of ideal types.

Three results stand out from among the implications of the model. First, in contrast to some of the recent literature, I find that unbiased mediators are generally more capable to help settle conflicts than biased mediators.[6] An increase in bias, however, does not always imply that a settlement becomes less likely. In particular this is true for "hard" cases, in which no settlements that would be acceptable to both sides exist when negotiations begin, but a compromise is possible as long as disputants change their expectations in response to the mediation effort. In these cases, unbiased mediators are most effective, but strongly biased mediators are more effective than weakly biased mediators in bringing about a settlement.

Second, "hard" cases also tend to be settled if the intervening mediator is not averse to allowing the dispute to continue. The effect of such a mediator is not, however, the same across the entire range of cases. Relatively tractable, "easy" cases have a smaller chance of being settled if the mediator does not mind when bargaining fails. I also find that unbiased mediators choose to stay on the sidelines if they do not care to avoid continued conflict, which means that mediators that do not care to avoid continued conflict tend to be biased (and hence generally less effective) as well.

Third, I find that mediator selection does affect the type of mediator we observe during conflict management itself. For example, I find that biased mediators choose not to engage with the most intractable disputes, in which a range of potential agreements neither exists, nor will come into existence through mediation (given a biased mediator). Unbiased mediators, on the other hand, do not condition their choice to offer mediation on a dispute's apparent tractability, but on how expensive it is to provide mediation, relative to the disutility derived from refusing to make

[6] Touval and Zartman (1989); Kydd (2003), and Smith and Stam (2003) argue in favor of biased mediation. Rauchhaus (2006) finds that partiality detracts from mediator effectiveness.

an offer and watching the dispute continue. By implication, if the fixed costs of mediation are high, we are more likely to see biased mediators stepping in to help resolve conflict.

2 The structure of the game

Consider a game involving a mediator M and two actors $\{1, 2\}$ who are bargaining over some good with a total value normalized to 1. In this game, nature first chooses some distribution of the good $\{\theta, 1-\theta\}$, which will be realized if the disputants reject a settlement proposed by the mediator. Let θ signify the share of the good received by player 1 if the dispute is not settled, without loss of generality. Neither of the disputants know the true value of θ, but each disputant i believes *a priori* that θ is distributed according to some probability density function $f_i(\theta)$ with support in $[0, 1]$, expectation p_i, and corresponding cumulative density function $F_i(\theta)$. If no settlement is reached, each disputant i receives a share of the good as determined by θ, minus his cost of conflict $v_i > 0$.

While the mediator M knows $f_i(\theta)$ and v_i for $i \in \{1, 2\}$, the disputants do not know each other's costs of conflict and beliefs about the distribution of θ. If they did, their expectations about the outcome of the dispute in the absence of a settlement would converge; they would not be able to rationally "agree to disagree" (Aumann, 1976). The intuition behind this fact is that each disputant's belief about what he will be able to secure for himself in the absence of a settlement reflects all knowledge the disputant holds in relation to the conflict. If that knowledge was shared by the disputants, it would mean that both disputants' expectations would draw on the same body of information and hence converge. (Note that expectations would not necessarily converge on θ, since the true outcome absent a settlement can be a function of factors neither disputant is aware of.) While this would reduce the potential for conflict, it is not in the strategic interest of the disputants to share their honest assessments of the situation. Each disputant would prefer to inflate his expected gains, if he knew that his statement was going to be believed, but this in turn prevents any statement from being credible.

In order to model the informational advantage of the mediator, assume that (a) she knows $f_i(\theta)$ and v_i for $i \in \{1, 2\}$ before her decision to offer mediation or not, (b) she observes the true θ if she does mediate the dispute (i.e. offers to mediate and is invited by the disputants). This captures the idea that mediators have a sense of where disputants stand before getting involved, and that once mediators do get involved, they find out where the conflict is headed in the absence of a settlement. This seems to approximate real-world processes better than other possible sets of assumptions. For instance, if we assumed that mediators did not know the disputants' costs and beliefs about θ, the model would have the unrealistic implication that mediators fail to differentiate between conflict situations in their decision to offer mediation or not (specifically if we assumed that the mediator held uninformed, uniform prior beliefs about $f_i(\theta)$ and v_i, then all types of mediators would always choose to offer their services). If we conversely assumed that mediators know not only disputants' costs and expected payoffs but also θ prior to their decision to offer mediation, they would perfectly select only into those conflicts in which a settlement can be achieved with certainty, again an unrealistic assumption. Since the mediator does not know θ with certainty prior to becoming involved in the conflict, her initial belief is that θ is distributed according to a probability density function $f_M(\theta)$ with support in $[0, 1]$, expectation p_M, and cumulative density function $F_M(\theta)$.

There are a number of reasons why mediators can have privileged access to dispute-relevant information: They can have fewer cognitive or social biases in evaluating available information; disputants sometimes disclose facts to mediators they would not disclose directly to one another; mediators can sometimes provide insights on contextual factors the disputants may be unaware of, such as international contingency plans for the dispute at hand; and mediators may have access to military intelligence that is unavailable to the disputants. The latter is arguably what happened in the 1990 Indo-Pakistani crisis, which subsided only when Robert M. Gates, George Bush's Deputy National Security Advisor, was dispatched to the region.[7]. Since India's Brasstacks exercises of January 1987, there had not been such a "real possibility of war" (Hagerty, 1995: 92). The Bush administration was concerned that both Pakistan and India overestimated their expected gains from an escalated conflict. As Donald Kerr, now Director of the U.S. National Reconnaissance Office, told Seymour Hersh: "The intelligence community believed that without some intervention the two parties could miscalculate—and miscalculation could lead to a nuclear exchange" (Hersh, 1993: 64–6). Gates met with Pakistani President Ghulam Ishaq Khan and Chief of Army Staff (COAS) General Mirza Aslam Beg on May

[7]Tensions between India and Pakistan escalated when Indian police killed a number of demonstrators that were in violation of a curfew in Kashmir on January 20. On March 13, Pakistani Prime Minister Benazir Bhutto threatened a "thousand-year war," and Indian Prime Minister V. P. Singh suggested in turn that Pakistan would not last "a thousand hours of war" (Hagerty, 1995: 98–9)

20, and conveyed that "Washington had thoroughly war-gamed the Indo-Pakistani confrontation, and Pakistan was the loser in every scenario" (Hersh, 1993: 67–8). In New Delhi, on the other hand, Gates told Singh and Indian COAS V. N. Sharma, among others, that India would win, but was underestimating the costs of a nuclear confrontation—in the run-up to the Gates mission, the U.S. had gathered intelligence that indicated Pakistan had moved materiel from an alleged nuclear holding site to an air force base in Baluchistan, a province in West Pakistan (Hagerty, 1995: 102). In early June, the crisis was defused when India withdrew the armored units it had deployed to the Mahajan area near the border to Pakistan earlier that year.[8]

The sequence of moves in the game is as follows: As mentioned before, nature first selects a value of θ. Second, the mediator decides whether to offer mediation or not. If no mediation is offered, the game ends without a settlement. Third, if mediation is offered, each disputant decides whether to accept the offer or not. If either one or both of the disputants reject the offer, the game ends without a settlement. Fourth, if both disputants accept the offer of mediation, the mediator observes θ and proposes a division $\{m, 1-m\}$ of the good at stake, with m being the share that would go to player 1. Fifth, each disputant accepts or rejects the mediator's proposal, but both need to accept the proposal in order for it to be implemented. Figure 1 depicts the sequence of moves graphically. Payoffs are shown in the following order: Disputant 1, disputant 2, and the mediator.

The disputants' utilities follow straightforwardly from what was discussed above. For player 1 we have

$$u_1 = \begin{cases} m & \text{if a settlement is reached;} \\ \theta - v_1 & \text{if no settlement is reached.} \end{cases}$$

Similarly,

$$u_2 = \begin{cases} 1-m & \text{if a settlement is reached;} \\ 1-\theta - v_2 & \text{if no settlement is reached.} \end{cases}$$

[8]Hagerty challenges the version of events presented here and argues that "existential nuclear deterrence was the most important factor in South Asia's 'non-war' of 1990" (112). According to Hagerty, "Islamabad and New Delhi quietly used the Gates intervention as an excuse to deescalate the crisis" (106); "both sides were anxious to back away from the brink of war, and ... Gates provided them with a mechanism for doing so without appearing weak" (117).

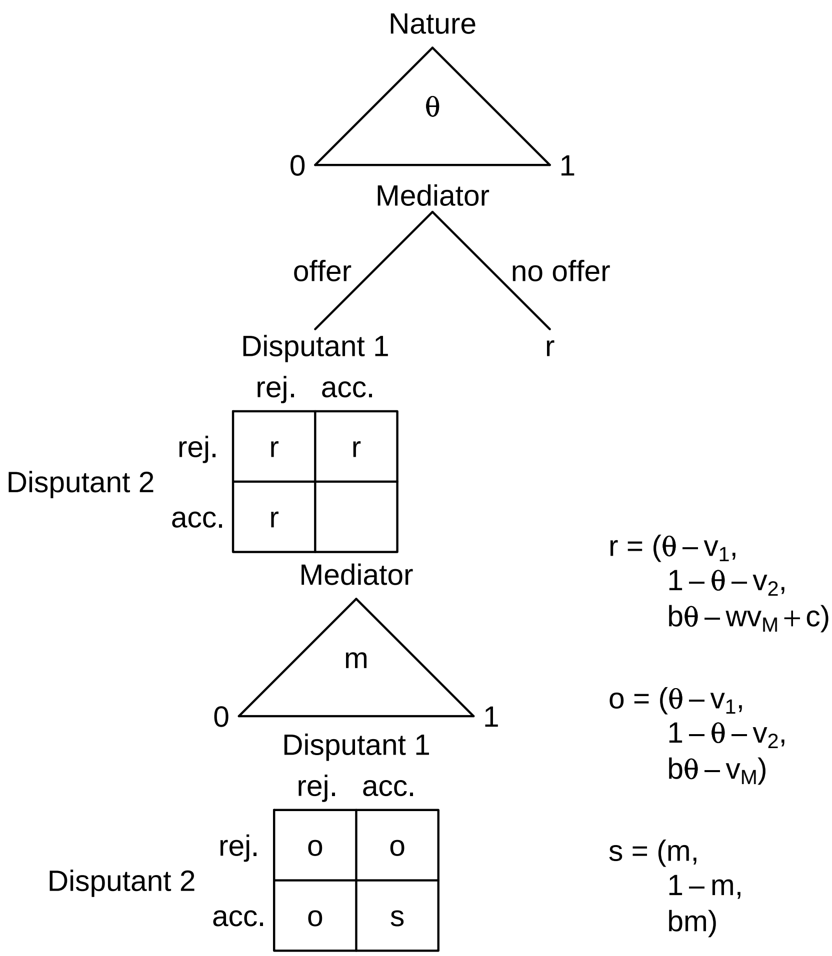

Figure 1: Game tree

For the mediator's utility we have

$$u_M = \begin{cases} bm & \text{if mediation occurs and a settlement is reached;} \\ b\theta - v_M & \text{if mediation occurs and no settlement is reached;} \\ b\theta - wv_M + c & \text{if no mediation occurs;} \end{cases}$$

where $v_M > 0$ is a cost associated with bargaining failure, $w \in [0, 1)$ is a weight used to scale the cost of bargaining failure in cases where the mediator does not become a party to the bargaining process, $c > 0$ is a fixed cost of providing mediation (or, in turn, a fixed benefit derived from not providing mediation), and b captures the mediator's bias in favor of one side of the dispute (expressed in terms of a weight that determines the amount of utility M derives from player 1's share of the stakes). Without loss of generality, assume that $b \geq 0$, meaning that the mediator is either unbiased or biased towards player 1.

This description of the mediator's utility captures important aspects of a third party's interests. First, mediators likely care about the distributional consequences of their work; they have a preference over what ought to happen with the good at stake, sometimes favoring one side over the other. This is captured by bm and $b\theta$, respectively. Second, mediators suffer when a settlement fails to materialize, similar to although conceivably to a lesser extent than disputants. I assume that this loss of utility hits the mediator both when she is involved in trying to reach a settlement and when she is not, although always to a smaller extent in the latter case than in the former. This component of the mediator's utility is characterized by v_M and wv_M respectively. Finally mediation consumes some of the intermediary's resources, as specified by parameter c.

Although this simplified model of dispute bargaining does not explicitly allow for a settlement to be reached in the absence of mediation, this is not a strict assumption. The distribution described by θ can be realized in any number of ways, including all-out war, a settlement negotiated directly between the two adversaries, or later mediation attempts. The assumption made is simply that no matter which way θ is realized, each player i suffers a loss of size v_i for failure to reach a settlement in the course of the game described here. Conversely this could be interpreted as an assumption that mediation is the least costly of any alternative ways of ending the dispute. The game does, in any case, not assume that mediation is the only way of ending the conflict.

In the following I will solve for all pure-strategy subgame-perfect equilibria of this

Symbol	Meaning
M	Mediator
$i \in \{1, 2\}$	Disputant
$\theta \in [0, 1]$	Player 1's share of the stakes if conflict continues
$1 - \theta$	Player 2's share of the stakes if conflict continues
$f_j(\theta)$ for $j \in \{1, 2, M\}$	Probability density function according to which Player j believes θ is distributed
$F_j(\theta)$ for $j \in \{1, 2, M\}$	Corresponding cumulative density function for player j
$p_j \in [0, 1]$ for $j \in \{1, 2, M\}$	Player j's prior expected value for θ
$v_j > 0$ for $j \in \{1, 2, M\}$	Cost incurred by player j if no settlement is reached
$b \geq 0$	Mediator's bias (i.e. a weight that captures the extent to which the mediator derives utility from an outcome that favors player 1)
$w \in [0, 1)$	Weight used to scale down M's cost of bargaining failure v_M if she does not mediate
$c > 0$	Fixed cost of mediation
θ_0	Threshold value below and above which the mediator takes different actions
$p_i^u \in [0, \theta_0]$	Disputant i's updated expected value for $\theta \in [0, \theta_0]$
$p_i^{u'} \in [\theta_0, 1]$	Disputant i's updated expected value for $\theta \in [\theta_0, 1]$
$p'_M \in [0, \theta_0]$	Mediator's prior expected value for $\theta \in [0, \theta_0]$

Table 1: Notation

game, starting at the end of the game tree. Pareto dominance is used throughout as a tool of equilibrium selection (i.e., at each node I reject any Nash equilibrium for which there exists some other Nash equilibrium which leaves no player worse off and at least one player better off).

3 The disputants' choice to accept or reject the mediator's proposal

Consider first the disputants' decision to accept or reject an agreement m proposed by the mediator. Let p_i^u denote player i's updated expectation about θ, given the mediator's proposal m. If m conveys no information about the true value of θ, we have $p_i^u = p_i$. Player 1 will accept proposal m if it leaves him at least as well off as rejecting the proposal, that is he will accept m if $p_1^u - v_1 \leq m$ and reject the proposal otherwise in the pareto dominant equilibrium. Player 2 will correspondingly accept m if $p_2^u + v_2 \geq m$ and reject it otherwise.

We find additional equilibria by changing the above strategy profile so that 1 accepts m only if $p_1^u - v_1 < m$ and rejects all other proposals, or 2 accepts m only if $p_2^u + v_2 > m$ and rejects all other m, or both, but all of these equilibria are pareto inferior. Mostly they generate the same payoffs as the pareto dominant equilibrium, but they leave 2 worse off and 1's payoff unchanged if 1's equilibrium action is to reject $m = p_1^u - v_1$, and leave 1 worse off while 2's payoff stays the same if 2 rejects $m = p_2^u + v_2$ in equilibrium. This suggests that switching to what I have suggested is the pareto dominant equilibrium does indeed constitute a pareto move.

The strategy profile by which both disputants reject all proposals is also in equilibrium (since neither actor has an incentive to deviate unilaterally), but it is pareto dominated by the equilibrium described above. In particular, rejecting all m is strictly pareto inferior in cases where $p_1^u - v_1 < p_2^u + v_2$ and $m \in [p_1^u - v_1, p_2^u + v_2]$ (and no different in terms of expected payoffs in all other cases).[9]

Given the pareto dominant equilibrium it must be the case that $p_1^u - v_1 \leq m \leq p_2^u + v_2$ for any m that will be accepted by both disputants. One can think of the interval $[p_1^u - v_1, p_2^u + v_2]$ as a bargaining space containing any acceptable agreements

[9] Note also that one could think to construct a set of equilibria in which one disputant plays any strategy (accept, reject, or a mixture of both) for m below both reservation values $p_2^u + v_2$ and $p_1^u - v_1$ while the other disputant rejects all m below his reservation point. This, however, is not an acceptable set of equilibria, because it requires one disputant to know the other's reservation value, which is inconsistent with my earlier assumption that disputant i does not know $f_j(\theta)$ for $j \neq i$.

m, with disputant 1's reservation value forming the lower bound and 2's reservation value at the upper bound. In the absence of such a bargaining space the dispute will not be settled.

Figure 2 shows how disputants' cost of continued conflict v_1 and v_2 can help generate such a bargaining space. If refusing a settlement was close to costless, the minimum share of the stakes that disputant 1 requires in order to agree to a settlement (i.e. $p_1 - v_1$) would be larger than the maximum share the other disputant would be willing to forgo (i.e. $p_2 + v_2$), and hence no settlement would be possible. But with v_1 and v_2 as substantial as they are in this example, a bargaining range of acceptable proposals opens up.

4 The mediator's proposal choice

4.1 Biased mediators

The mediator chooses m^* conditional on the disputants' equilibrium strategies at the final node. Consider first the case in which $b > 0$, which implies that the mediator's utility is strictly increasing in m and θ. In other words, if both m' and m'' are acceptable proposals and $m' > m''$, the mediator will optimally propose m'. It follows that any acceptable proposal will be located at the upper bound of the bargaining range, i.e. $m^* = p_2^u + v_2$.

4.1.1 Acceptable proposals: A complete pooling equilibrium

If no updating occurs in equilibrium, any acceptable proposal m^* will equal $p_2 + v_2$. Two conditions must be met in order for this proposal to be in equilibrium. First, we need $p_1 - v_1 \leq p_2 + v_2$ in order for the proposal to be acceptable to both disputants 1 and 2. The conflict must be tractable in the sense that a bargaining space already exists *ex ante*. We can think of such cases as "easy" cases, as opposed to "hard" cases in which no bargaining space exists prior to mediation. Figure 3 provides an example of an "easy" case, with the disputants' beliefs $f_i(\theta)$ for $i \in 1, 2$ generating a range of acceptable agreements.

Second, the mediator must have an incentive to keep negotiations from failing, no matter how much of an upper hand she believes the disputant she favors has. Her payoff from settlement m^*, which is $b(p_2 + v_2)$, must be at least as large as her payoff from bargaining failure $b\theta - v_M$, even if the mediator has observed a θ as high

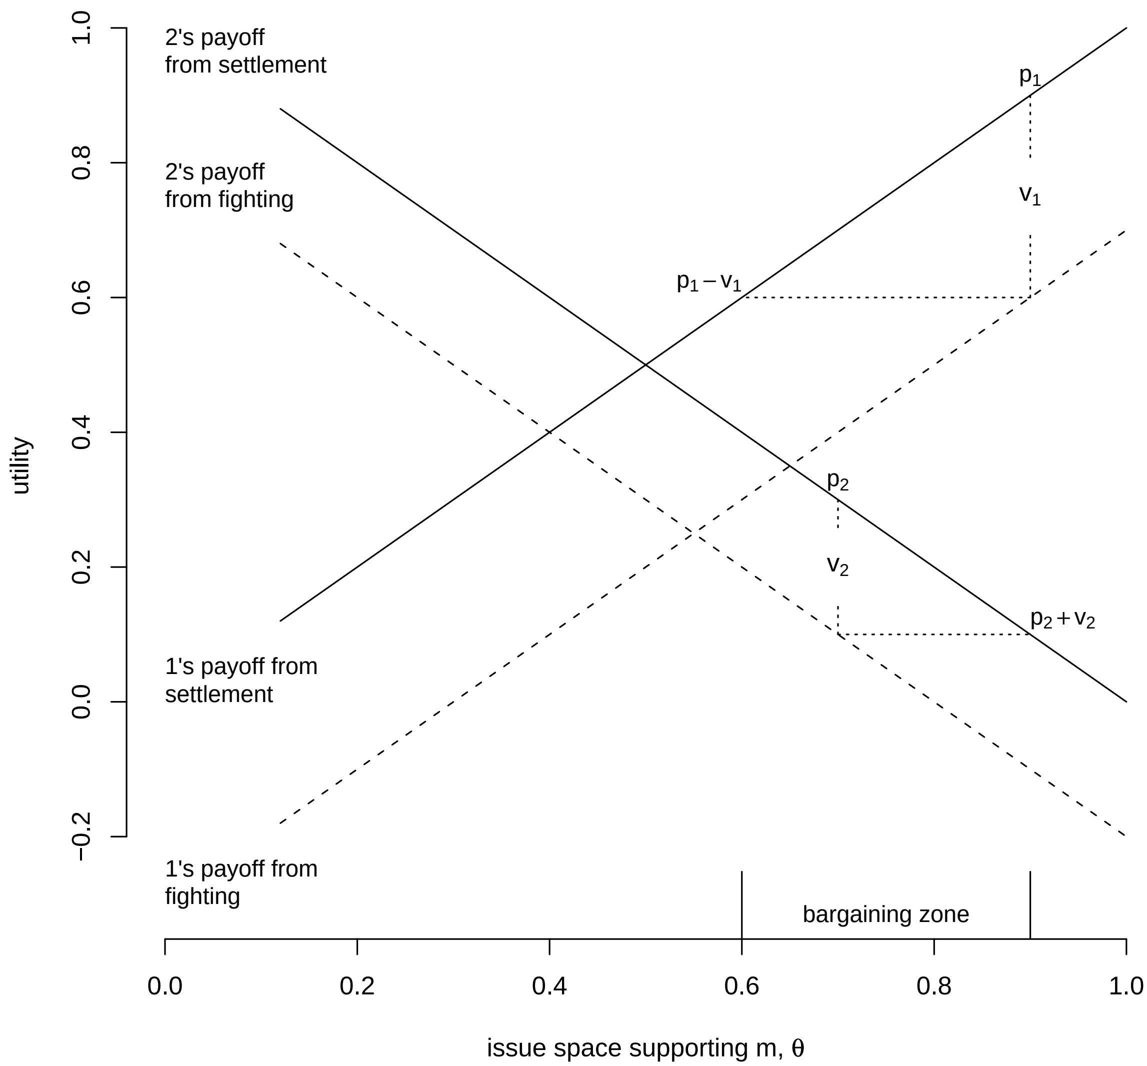

Figure 2: An example of a bargaining space

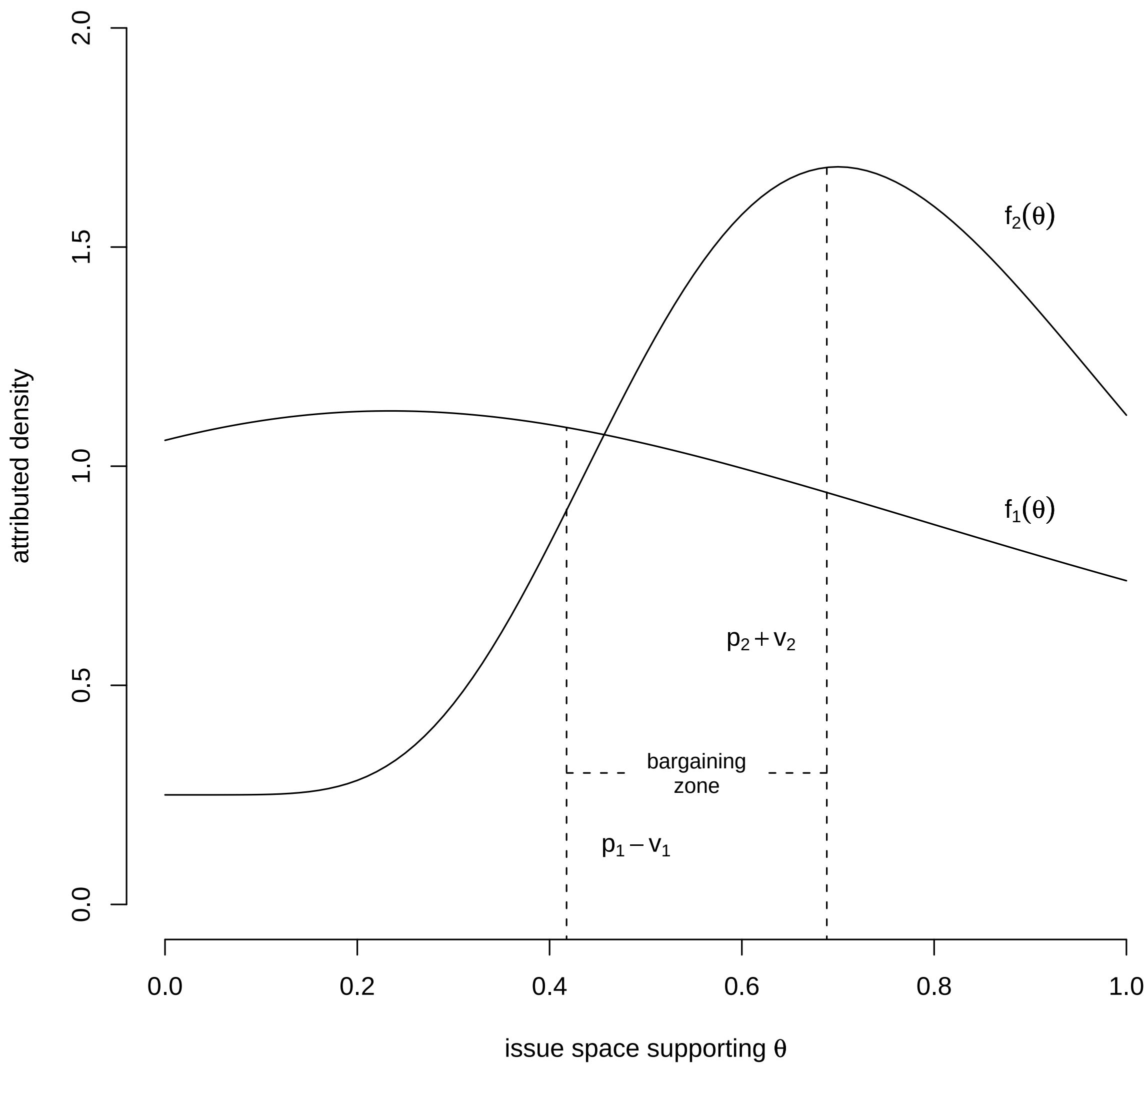

Figure 3: Expectations in an "easy" case

as 1. By simplifying we have that $p_2 + v_2 + \frac{v_M}{b} \geq 1$ must hold.[10]

Consider for example the Anglo-American mediation effort between Italy and Yugoslavia over Trieste in 1954. British and American negotiators were certainly biased toward Italy, a Cold War ally for whom the UK and the U.S. had administered the northern part of the disputed territory around Trieste, but even more importantly they had little interest in the dispute remaining unresolved after Tito had broken with Stalin. Finally, mediation did not need to change beliefs on either side in order to be effective, as both Italy and Yugoslavia already understood that the post-World War II demarcation line separating Anglo-American forces in Zone A from Yugoslav-controlled Zone B constituted an acceptable bargaining solution, despite the fact that Yugoslav delegates publicly protested American and British intentions to hand over the northern Trieste area to Italy (Zartman and Berman, 1982: 68). As the British negotiator Geoffrey Harrison put it: "They both claimed to have rights to more territory. They both knew they could not get it. They both realized that the best they could hope to get was Zone A and Zone B, respectively" (Campbell, 1976: 72–3).

4.1.2 Acceptable proposals: Partially separating equilibria

If disputants update their expectations about θ in equilibrium, any acceptable proposal $m^* = p_2^u + v_2$. Since $b > 0$, the mediator cannot credibly relay the exact value of θ, because if she was believed, her bias would drive her to inflate the value of θ. The mediator can, however, credibly separate some values of θ from others by proposing an agreement that will be accepted for θ at or below a certain threshold and making a proposal that will not be accepted for θ above this threshold (but not vice versa, because in that case the mediator would have an incentive to pretend that she observed a high θ when in fact she did not.) In this partially separating equilibrium, the mediator proposes $m^* = p_2^u + v_2$ for any θ she observes at or below the threshold $\theta_0 = p_2^u + v_2 + \frac{v_M}{b}$, and $m^* > p_2^u + v_2$ (which will not be accepted) for any θ above the threshold.[11]

[10] One could think that there exists an equilibrium in which M makes an unacceptable proposal if and only if $p_2 + v_2 + \frac{v_M}{b}$ is exactly equal to 1 and she observes $\theta = 1$, since this leaves her indifferent between a settlement and continued conflict. If the mediator was to take this particular action upon observing $\theta = 1$, however, disputants would know that the true θ is equal to 1 and update correspondingly, which in turn gives the mediator an incentive to take the action even if $\theta \neq 1$, in contradiction to the proposed equilibrium.

[11] A second partially separating equilibrium exists in which the mediator proposes $m^* = p_2^u + v_2$ for any θ she observes *strictly* below the threshold θ_0, and $m^* > p_2^u + v_2$ for any θ at or above the

Three conditions must hold in order for this equilibrium to exist, and since M prefers $m' > m''$ if both m' and m'' are acceptable proposals, θ_0 is always the largest of any candidate threshold values for which these conditions are true. First, it must be true that $p_2^u + v_2 + \frac{v_M}{b} < 1$, because otherwise the mediator would have an incentive to propose m^* no matter the value of θ, as in section 4.1.1 in which no updating occurred, thus ridding m^* of any informational content about θ and making it impossible for disputants to update. Second, $f_i(\theta)$ and v_i for $i \in \{1, 2\}$ must be such that $p_1^u - v_1 \leq p_2^u + v_2$ exists. Both disputants have a prior belief about how θ is distributed over the interval from 0 to the threshold value θ_0, and if no bargaining space exists within this interval, the partially separating equilibrium does not exist. Third, it must also be true that $f_1(\theta)$ and $f_2(\theta)$ are such that no bargaining space exists over the interval from θ_0 to 1. (This is always true if the previous condition holds and the disputants' beliefs fail to generate a bargaining space over the entire domain of θ from 0 to 1.) Otherwise the mediator would be able to put an acceptable proposal on the table for values of θ above the threshold, but if such a proposal was acceptable, the mediator would advance that same proposal even if θ was actually below the threshold.

Figure 4 shows an example of a "hard" case, in which no bargaining zone exists prior to mediation. When the mediator observes $\theta \leq \theta_0$ in this example, she proposes $p_2^u + v_2$, disputants update their expectations accordingly, and a settlement is reached. Here all three of the conditions discussed above are met.

4.1.3 Getting rejected: Other equilibria

No other equilibrium in which proposals are accepted in pure strategies (and $b > 0$) exists. Three additional features of this game should be noted: First, if the complete pooling equilibrium described in section 4.1.1 exists, no other pure strategy equilibrium of any kind can be found. Second, if the partially separating equilibrium exists, complete pooling equilibria in which the mediator proposes the same $m^* \in [0, 1]$ no matter which value she observed for θ exist as well (the mediator in that case does not have an incentive to deviate unilaterally, given that disputants will not update and the agreement will be rejected no matter the proposal), but the latter equilibria are pareto dominated by the former. For values of θ at or above θ_0,

threshold. This equilibrium is pareto dominated for all $p_1^u - v_1 < p_2^u + v_2$, because in that case an offer of $p_2^u + v_2$ leaves 1 strictly better and no other player worse off. For the special case of $p_1^u - v_1 = p_2^u + v_2$, however, it is not pareto dominated.

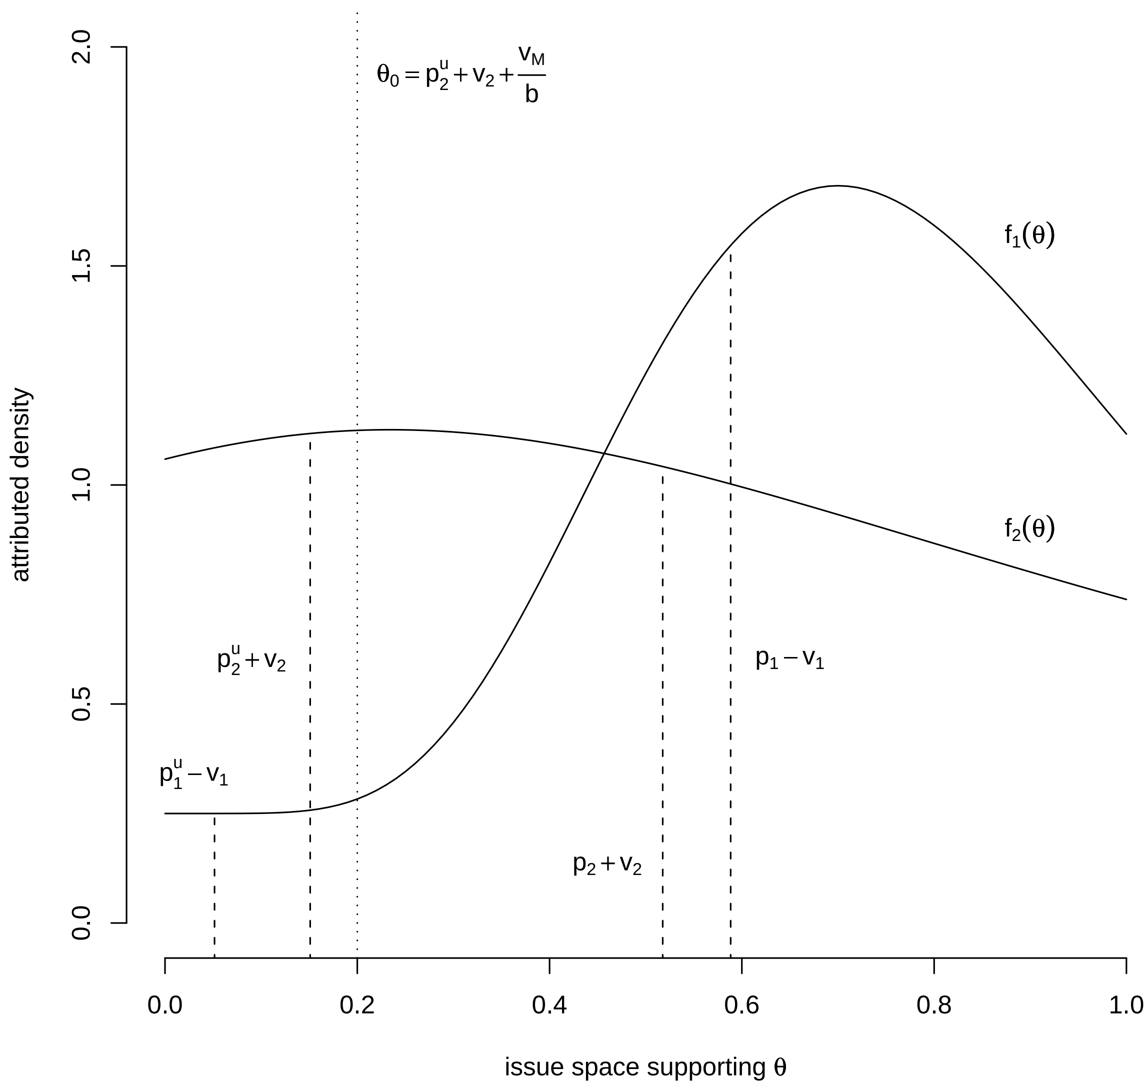

Figure 4: Expectations in a "hard" case

either type of equilibrium leads to the same payoff for the mediator, but for $\theta < \theta_0$ the partially separating equilibrium leaves the mediator strictly better off. Finally, if neither the complete pooling nor the partially separating equilibrium described in the previous subsections exist, then the only pure strategy equilibria that exist are complete pooling equilibria which fail to lead to an agreement (i.e., the mediator proposes $m^* \in [0, 1]$ regardless of θ and no updating occurs).

4.2 Unbiased mediators

Consider now the case where $b = 0$ (note that some of the conditions specified above are only defined for $b \neq 0$). The mediator's payoff then simplifies to 0 if an agreement is reached and $-v_M$ otherwise, which means that the mediator will want to propose any agreement that will be accepted. By implication, the mediator does not have an incentive to deviate from proposing $m^* = \theta$, the disputants recognize this, update their expectations so that $p_1^u = p_2^u = m^*$, and accept the proposal. This completely separating equilibrium always exists when $b = 0$.

Two types of completely pooling equilibrium can also be found. First, if $p_1 - v_1 \leq p_2 + v_2$, equilibria exist in which the mediator proposes the same $m^* \in [p_1 - v_1, p_2 + v_2]$ regardless of θ. Here a bargaining space in the disputants' beliefs already exists *ex ante*, and the unbiased mediator makes an acceptable proposal that falls within this bargaining range. Second, if $p_1 - v_1 > p_2 + v_2$, we have complete pooling equilibria in which the mediator proposes some $m^* \in [0, 1]$ regardless of θ, no updating occurs, and the proposal is rejected. The mediator has no incentive to deviate, since the disputants would believe that any $m' \neq m^*$ contained as little information about the true value of θ as m^*. This last group of equilibria, however, is pareto dominated by the separating equilibrium in which $m^* = \theta$. An unbiased mediator will always be better off in an equilibrium in which she conveys θ truthfully and reaches an agreement than in an equilibrium in which disputants fail to adjust their expectations.

5 Implementing mediation: A summary

When will disputants accept an agreement proposed by the mediator? There are, broadly, three types of cases in which a settlement can be accomplished. First, the dispute is settled if we have an unbiased mediator, no matter the value of any other parameters in the model and regardless of whether a bargaining space already existed prior to the mediation effort or not. One can think of cases in which reservation values fail to overlap *ex ante* as "hard" cases, and cases in which a zone of agreement already exists from the start as "easy." An impartial mediator will be able to help settle both the difficult and the easy disputes, and impartiality is all she requires in doing so.

Second, a biased mediator can help settle the "easy" cases if the mediator's cost associated with failing to reach an agreement is sufficiently large relative to the magnitude of her bias. Specifically, the ratio of the cost of mediation failure over the mediator's bias has to be at least as large as the distance between the total value of the good at stake and the maximum which the disputant not favored by the mediator is willing to forego at the outset.

Third, a biased mediator can also facilitate the settlement of "hard" cases, but her chances of doing so are limited. First, her bias must be large relative to the cost associated with bargaining failure, which is the opposite of what must be true in "easy" cases. The ratio of cost over bias has to be smaller than the distance between the good's total value and the share that an acceptable proposal would give to the disputant not favored by the mediator. Second, if the mediator conveys to the disputants that continued conflict will lead to a situation in which the mediator's favored disputant is left with relatively little or possibly nothing at all of the stakes, it must be true that when disputants change their expectations in the face of this information, they find that they can agree on a settlement. Third, in order to make the mediator's statement credible, it must be true that disputants will not find themselves able to agree on a settlement if the mediator's proposal suggests that her favored disputant will receive a relatively large share of the stakes as a result of continued conflict. This captures an idea from the literature on biased advisers and intermediaries: An individual who is partial toward the recipient of his message will be able to effectively convey bad news only, because good news need not have anything to do with the truth in order for the individual to be willing to share them (Calvert, 1985).

6 The disputants' procedural choice

Each disputant's decision to accept or reject an offer of mediation does not mirror whether mediation would lead to a settlement or not. Disputants instead weigh the distributional consequences of any potential mediation outcome against their expected payoffs from declining the mediator's offer. Consider first the case in which mediation fails to lead to a settlement. Then both disputants are indifferent between allowing and refusing mediation, since either way they will receive a payoff of $\theta - v_i$ for $i \in \{1, 2\}$. Second, if mediation does lead to a settlement m^*, we know from section 3 that $p_1^u - v_1 \leq m^* \leq p_2^u + v_2$ (where $p_i^u = p_i$ for $i \in \{1, 2\}$ if no updating occurs), so disputants' expected utilities will be at least as large under the agreement as in the absence of mediation. A strategy profile by which both disputants always accept mediation is hence in equilibrium.

We can further show that this equilibrium pareto dominates all other equilibria. There are two other sets of pure strategy equilibria. First, if the mediator is biased ($b \neq 0$), we know from section 4.1 that any acceptable m^* will equal $p_2^u + v_2$ and hence be equivalent to player 2's expected payoff from refusing mediation. In turn player 2 has no incentive to deviate from always declining mediation, and player 1 either accepts or rejects the offer of mediation in these equilibria. However, player 1 would be better off and player 2 no worse off if they both accepted mediation, regardless of any other parameters, which means that this equilibrium is pareto dominated.[12] Second, a pure strategy equilibrium exists in which both disputants always reject the offer of mediation (regardless of b), since neither disputant can improve his payoff by deviating unilaterally. The disputants, however, receive higher utilities in the equilibrium in which they jointly accept mediation. If the mediator is unbiased, both disputants are strictly better off with mediation (gaining v_1 and v_2 in expectation respectively), and if the mediator is biased, player 1 improves his payoff by the same reasoning as above. Either way, the equilibrium in which disputants always reject mediation is pareto dominated.

There are two important reasons for why disputants always accept mediation in this model. First, the model assumes that all fixed costs of mediation are absorbed by the mediator, and that there is no cost for the disputants associated with delaying the outcome that will transpire in the absence of a mediated settlement: Not agreeing to mediation in the first place yields the same payoff as accepting mediation and failing to reach an agreement later. Second, there is no more than one mediation offer on the table at a time. Disputants are not deciding who to choose from among

[12] In some cases the disputants' prior beliefs are such that biased mediation cannot possibly reach a settlement, and in those instances player 1 is not strictly better off with mediation. The disputants, however, do not know each other's prior beliefs, which means that they operate under the assumption that there is some chance that biased mediation will lead to an agreement, regardless of any other parameters.

a group of potential mediators. While these are simplifying assumptions, they imply that as long as disputants incur no direct costs for participating in mediation (and as long as there are no competing offers, in which case some offers would be rejected be necessity), mediation will be allowed to proceed. The model also implies that mediation does not have to be costly for disputants in order to lead to a settlement.

7 The offer to mediate

There are two reasons for why a mediator can have an incentive not to offer mediation in the model. First, there is a fixed cost c associated with providing mediation services. Second, the mediator's negative utility of continued conflict v_M is scaled down by a factor of w if she does not become involved in the dispute. Given the costs associated with mediation, the mediator will not offer her services if an agreement cannot be reached. If $b > 0$, this means that no mediation will be offered if (a) either $p_2 + v_2 \geq p_1 - v_1$ or $p_2 + v_2 + \frac{v_M}{b} \geq 1$ do not hold, and (b) there does not exist a threshold θ_0 such that $p_2^u + v_2 \geq p_1^u - v_1$ and $p_2^u + v_2 + \frac{v_M}{b} < 1$, where p_i^u denotes i's expectation about θ over the interval $[0, \theta_0]$.

But even if an agreement is possible, mediation is not inevitable. If $b = 0$, the mediator's payoff from entering the dispute and helping to reach an agreement is 0, and the payoff from staying out of the dispute is $wv_M + c$, which means that we need $wv_M \geq c$ in order for mediation to occur. (Not offering mediation when $wv_M = c$ is in equilibrium, but it is pareto dominated by the equilibrium in which the mediator enters, because M is no worse off in doing so while the disputants' payoffs improve by a margin of v_1 and v_2 respectively.)

If the mediator is biased, there are two types of agreements that can be reached. First, there are "easy" cases, in which a bargaining space exists from the start and $p_2 + v_2 + \frac{v_M}{b} \geq 1$. A mediation offer is in equilibrium if M's payoff from entering the dispute, which is $b(p_2 + v_2)$, is at least as large as her expected payoff from refusing to mediate, which is $bp_M - wv_M + c$. Recall that prior to observing θ, the mediator expects θ to take on the value of p_M. By simplifying we have that

$$p_2 + v_2 + \frac{wv_M - c}{b} \geq p_M \quad (1)$$

must hold.[13] Note that $p_2 + v_2 + \frac{wv_M - c}{b}$ is increasing in v_M and w, and decreasing

[13]There also exists an equilibrium in which M does not make an offer when (1) holds with strict

in c. All else equal, the greater the mediator's loss of utility if no agreement is reached, and the greater the extent to which the mediator incurs this loss even if she doesn't become involved, the more likely she is to offer mediation. She is the less likely to offer mediation, the greater the fixed procedural costs of mediation.

The relative sizes of w, v_M, and c also determine the marginal effect of a shift in bias on the likelihood of mediation. If $wv_M > c$, an increase in bias makes mediation less probable, but for $wv_M < c$ a decrease in bias has that effect. If the costs of the mediation process and bargaining failure are such that an unbiased mediator would (not) offer her services, then greater bias means mediation becomes less (more) likely.

Second, I consider "hard" cases in which no bargaining space exists *ex ante*, but an agreement is possible because there exists a threshold $\theta_0 = p_2^u + v_2 + \frac{v_M}{b} < 1$ such that $p_2^u + v_2 \geq p_1^u - v_1$. In order for M to prefer offering mediation, we require that her payoff from doing so is at least as great as her payoff from remaining on the sidelines. Formally,

$$F_M(\theta_0)b(p_2^u + v_2) + (1 - F_M(\theta_0))\left(b\int_{\theta_0}^{1} \theta f_M(\theta)\,d\theta - v_M\right) \geq bp_M - wv_M + c$$

where $b(p_2^u + v_2)$ is the mediator's payoff from reaching a settlement, weighted by the probability $F_M(\theta_0)$ of θ falling below threshold θ_0; the cost of bargaining failure v_M is deducted from the product of the mediator's bias b and the integral denoting the expected value of θ above threshold θ_0 to yield the mediator's payoff from intervening but failing to reach a settlement, a payoff which is weighted by the probability $1 - F_M(\theta_0)$ of θ falling above threshold θ_0; and $bp_M - wv_M + c$ is the mediator's payoff from not becoming involved in the dispute in the first place. We can simplify to

$$F_M(\theta_0)\left(p_2^u + v_2 + \frac{v_M}{b} - p'_M\right) + \frac{(w-1)v_M - c}{b} \geq 0 \qquad (2)$$

where $p'_M = \int_0^{\theta_0} \theta f_M(\theta)\,d\theta$ is the mediator's prior expectation of θ in the interval $[0, \theta_0]$.[14] As in the "easy" cases, this equation is the more likely to hold (and

equality. If p_2+v_2 is strictly smaller than p_1-v_1, this equilibrium is pareto dominated, since offering mediation would generate a higher payoff for disputant 1 and leave all other payoffs unchanged. This other equilibrium is not pareto dominated for the special case of $p_2 + v_2 = p_1 - v_1$.

[14] By the same reasoning as in footnote 13, an equilibrium exists in which M does not make an

mediation is the more likely to be offered), the greater the extent w to which the mediator incurs a loss in utility from bargaining failure even if she does not offer to mediate and the smaller the fixed procedural costs c of mediation. A more subtle picture emerges with respect to the third party's bias b and cost of bargaining failure v_M. By taking the partial derivative of the LHS of equation (2) with respect to v_M and b we have

$$\frac{\partial}{\partial v_M} = \frac{F_M(\theta_0) + w - 1}{b} \tag{3}$$

and

$$\frac{\partial}{\partial b} = \frac{c - v_M(F_M(\theta_0) + w - 1)}{b^2} \tag{4}$$

which suggests two things. First, since equation (3) will be positive for $1 - F_M(\theta_0) < w$ and negative for $1 - F_M(\theta_0) > w$, it follows that an increase in the loss that the mediator incurs if no settlement is reached makes a mediation offer more likely as long as the probability that θ will be such that no settlement can be reached is sufficiently small (i.e. smaller than w). If the probability that no settlement will emerge is relatively high, however, a decrease in the mediator's loss of utility from bargaining failure will make mediation more likely. If chances for a settlement are dim, the more insensitive a mediator to whether she can make a settlement happen, the more likely it is that she will step in the fray. If a settlement is likely, the more a mediator cares about an agreement being reached, the more likely she is to offer her services.

Second, equation (4) will be positive if $1 - F_M(\theta_0) > w - \frac{c}{v_M}$ and negative if $1 - F_M(\theta_0) < w - \frac{c}{v_M}$, which means that the effect of bias on the likelihood of a mediation offer is conditional on M's expectation about how likely it is that a settlement can be reached. If there is little chance of a settlement (i.e. the expected probability of θ falling above threshold θ_0 is larger than $w - \frac{c}{v_M}$), then a more biased mediator is more likely to make an offer. If a settlement is relatively probable, a more biased mediator is less likely to become involved.

Finally, note that the marginal effects of changes in the exogenous parameters

offer when (2) holds with strict equality, and this equilibrium is not pareto dominated in the special case of $p_2^u + v_2 = p_1^u - v_1$.

c, w, v_M, and b are more pronounced for mediators with low bias, because of the bias terms in the denominator of equations (1), (3), and (4). If a mediator exhibits little bias to start with, changes in the mediator's cost and bias parameters have a relatively larger effect on the likelihood of mediation being offered, regardless of whether a dispute constitutes an "easy" or a "hard" case.

8 Initiating mediation: A summary

When do mediators become involved in a dispute? I have described a mediator's incentives in terms of four components: (1) her bias toward one of the disputants; (2) the fixed costs of providing mediation; (3) the cost she incurs if the dispute is not settled; and (4) the extent to which she incurs the latter cost even if she does not offer to mediate. Both (2) and (4) have an unambigious and straightforward impact on the likelihood that a mediator will provide mediation. If the cost of bargaining failure even in the absence of mediation is substantial, and if mediation is cheap to provide, the mediator is more likely to offer her assistance to the disputants.

An increase in the mediator's cost of bargaining failure, the third element of the incentives listed above, leads to a higher probability of mediation in all but the toughest cases. Only if no range of acceptable agreements exists at the outset (i.e. disputants must change their expectations about the consequences of continued conflict in the course of mediation), and the probability that mediation will not produce an agreement is particularly high, then a decrease in the third party's cost of bargaining failure is associated with a greater probability of mediation.

The effect of bias is similarly contingent. In "easy" cases (i.e. disputes that have reached a point where what is minimally acceptable to one disputant is also at least minimally acceptable to the other disputant, even before mediation) a more biased mediator is less likely to make a mediation offer to the disputants as long as the fixed costs of mediation are small relative to the loss the mediator suffers if no settlement is reached. If mediation is expensive to provide relative to the cost of failing to reach an agreement, an increase in bias means a better chance of mediation being offered. In "hard" cases, the effect of bias on mediation onset also depends on how likely it is that a range of acceptable agreements will emerge when mediation is offered: A dispute that has a particularly high (low) chance of remaining intractable attracts a more (less) biased mediator.

Biased mediator?	Does a bargaining space exist *ex ante*?	Can a bargaining space open up?	Do the mediator's returns outweigh her costs of mediation?	Case
$b = 0$			$wv_M \geq c$	1
			$wv_M < c$	2
$b > 0$	$p_2 + v_2 \geq p_1 - v_1$, and $p_2 + v_2 + \frac{v_M}{b} \geq 1$		$p_2 + v_2 + \frac{wv_M - c}{b} \geq p_M$	3
			$p_2 + v_2 + \frac{wv_M - c}{b} = p_M$, and $p_2 + v_2 = p_1 - v_1$	4
			$p_2 + v_2 + \frac{wv_M - c}{b} < p_M$	5
	$p_2 + v_2 < p_1 - v_1$, or $p_2 + v_2 + \frac{v_M}{b} < 1$	$p_2^u + v_2 + \frac{v_M}{b} < 1$, and $p_2^u + v_2 \geq p_1^u - v_1$, and $p_2^{u'} + v_2 < p_1^{u'} - v_1$	$F_M(\theta_0)\left(p_2^u + v_2 + \frac{v_M}{b} - p'_M\right) + \frac{(w-1)v_M - c}{b} \geq 0$	6
			$F_M(\theta_0)\left(p_2^u + v_2 + \frac{v_M}{b} - p'_M\right) + \frac{(w-1)v_M - c}{b} = 0$ and $p_2^u + v_2 = p_1^u - v_1$	7
			$F_M(\theta_0)\left(p_2^u + v_2 + \frac{v_M}{b} - p'_M\right) + \frac{(w-1)v_M - c}{b} < 0$	8
		$p_2^u + v_2 + \frac{v_M}{b} \geq 1$, or $p_2^u + v_2 < p_1^u - v_1$, or $p_2^{u'} + v_2 \geq p_1^{u'} - v_1$		9

Table 2: Conditions that define equilibrium actions and outcomes

Biased mediator?	Bargaining space exists *ex ante*?	Can a bargaining space open up?	Do the mediator's returns outweigh her costs of mediation?	Case
No			Yes	1
			No	2
Yes	Yes, and mediator has incentive to propose agreement in bargaining space		Yes	3
			Mediator indifferent; disputants indifferent between settlement and continued conflict	4
			No	5
	No, or mediator does not have incentive to propose agreement in bargaining space	Yes, and mediator wants acceptable proposal, and agreements best for disputant not favored by mediator are not acceptable	Yes	6
			Mediator indifferent; disputants indifferent between settlement and continued conflict	7
			No	8
		No, or mediator does not want acceptable proposal, or agreements best for disputant not favored by mediator are acceptable		9

Table 3: Conditions that define equilibrium actions and outcomes

Case	Mediation offered?	Observed θ	Mediator's proposal m^*	Disputant's posterior belief	Proposal accepted?	Payoffs for 1, 2, and M	Outcome
1	Yes	Any	θ	θ	Yes	θ $1-\theta$ 0	Settlement
2	No					$\theta - v_1$ $1-\theta+v_2$ $b\theta - wv_M + c$	Continued conflict
3	Yes	Any	p_2+v_2	p_i	Yes	p_2+v_2 $1-p_2+v_2$ $b(p_2+v_2)$	Settlement
4	No					$\theta - v_1$ $1-\theta+v_2$ $b\theta - wv_M + c$	Continued conflict
5	No					$\theta - v_1$ $1-\theta+v_2$ $b\theta - wv_M + c$	Continued conflict
6	Yes	$[0,\theta_0]$	$p_2^u+v_2$	p_i^u	Yes	$p_2^u+v_2$ $1-p_2^u+v_2$ $b(p_2^u+v_2)$	Settlement
		θ_0, and $p_2^u+v_2=p_1^u-v_1$	$(p_2^u+v_2,1]$	$p_i^{u'}$	No	$\theta - v_1$ $1-\theta+v_2$ $b\theta - v_M$	Continued conflict
		$(\theta_0,1]$	$(p_2^u+v_2,1]$	$p_i^{u'}$	No	$\theta - v_1$ $1-\theta+v_2$ $b\theta - v_M$	Continued conflict
7	No					$\theta - v_1$ $1-\theta+v_2$ $b\theta - wv_M + c$	Continued conflict
8	No					$\theta - v_1$ $1-\theta+v_2$ $b\theta - wv_M + c$	Continued conflict
9	No					$\theta - v_1$ $1-\theta+v_2$ $b\theta - wv_M + c$	Continued conflict

Table 4: Description of equilibrium path in pareto dominant, pure strategy, subgame-perfect equilibria

9 Conclusion

This paper presented a formal analysis of the strategic choices disputants and mediators make when they decide to engage in and carry out mediation. Tables 2 and 3 describe the conditions under which different scenarios occur in equilibrium, and table 4 shows the corresponding choices, payoffs, beliefs, and outcome that obtain on the equilibrium path.

In concluding I want to highlight three sets of findings. First, the model suggests that unbiased mediators outperform biased ones in terms of persuading disputants to accept a settlement, regardless of other factors and of prior expectations about the results of continued conflict. Furthermore if a biased mediator is involved, "easy" cases (i.e. disputes in which a zone of agreement exists at the outset of negotiations) are more likely to be resolved if the mediator's bias is not particularly severe. "Hard" cases, on the other hand, are more likely to be settled with the help of a strongly rather than a weakly partial third party, again given a biased mediator. Perhaps it is this discontinuity in the effect of mediator bias—with unbiased mediators best equipped to settle conflict, but seriously biased mediators better equipped than slightly biased ones to handle tough cases—that has helped spawn a continuing debate over its merits.

Second, the model speaks to the question of whether a mediator with a strong incentive to avoid the breakdown of negotiations is preferable to a mediator who is not seriously affected if the conflict continues without a settlement. The model presented in this paper suggests that, if the mediator in question is unbiased, a settlement is more likely if the mediator does not want bargaining to fail—not because it helps in the process of mediation itself, but because it increases the likelihood that the unbiased third party chooses to make an offer of mediation. If we consider a biased mediator, then a large cost linked to bargaining failure helps in the resolution of "easy" cases, but "hard" cases are better tackled by a mediator who cares little about whether or not a settlement is reached (and worries instead about what kind of settlement can transpire).

Third, the model does not only consider the effects of mediator characteristics on mediation once initiated, but also asks whether particular types of mediators are selected into particular types of conflict. In summary, two results should be noted. First, biased mediators steer clear of the toughest, most intractable conflicts (i.e. those where no bargaining space exists to start with, nor the possibility of a bargaining space opening up if the mediator was to propose an agreement cutting against her bias), while unbiased mediators are principally willing to engage with cases that can range from completely ripe for resolution to thoroughly intractable. Second, unbiased mediators select into cases by weighing the fixed costs of providing mediation against the disutility of staying out and seeing the dispute rage on. If mediation is sufficiently expensive in terms of fixed procedural costs, unbiased mediators shirk from becoming involved, but biased mediators do not necessarily. If, in other words, mediation is difficult to orchestrate for technical reasons, biased mediators are the ones who will make themselves available for assistance anyhow, with the consequence of comparatively higher rates of bargaining failure than if mediator characteristics were assigned randomly. This example suggests that the process of mediation initiation can complicate an assessment of the implementation stage of mediation in international disputes, and that these phenomena should be studied jointly.

References

Robert J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.

Kyle Beardsley and Holger Schmidt. (Almost) All a matter of supply: Examining the determinants of UN involvement and effectiveness in international crises, 1945-2002. In *Proceedings of the Annual Meeting of the American Political Science Association*, Philadelphia, 2006.

Jacob Bercovitch. Introduction: Thinking about mediation. In Jacob Bercovitch, editor, *Resolving International Conflicts: The Theory and Practice of Mediation*, pages 1–10. Boulder: Lynne Rienner, 1996.

Jacob Bercovitch and Paul Diehl. Conflict management of enduring rivalries: The frequency, timing, and short-term impact of mediation. *International Interactions*, 22(3):299–320, 1997.

Randall L. Calvert. The value of biased information: A rational choice model of political advice. *The Journal of Politics*, 47(2):530–555, 1985.

John C. Campbell, editor. *Successful Negotiation: Trieste 1954*. Princeton: Princeton University Press, 1976.

William J. Dixon. Third-party techniques for preventing conflict escalation and promoting peaceful settlement. *International Organization*, 50(4):653–681, 1996.

James D. Fearon. Rationalist explanations for war. *International Organization*, 49 (3):379–414, 1995.

Ronald J. Fisher. Pacific, impartial third-party intervention in international conflict: A review and an analysis. In John A. Vasquez, James T. Johnson, Sanford Jaffe, and Linda Stamato, editors, *Beyond Confrontation: Learning Conflict Resolution in the Post-Cold War Era*, pages 39–59. Ann Arbor: University of Michigan Press, 1995.

Devin T. Hagerty. Nuclear deterrence in South Asia: The 1990 Indo-Pakistani crisis. *International Security*, 20(3):79–114, 1995.

Seymour M. Hersh. On the nuclear edge. *The New Yorker*, pages 56–73, March 29 1993.

Andrew Kydd. Which side are you on? bias, credibility and mediation. *American Journal of Political Science*, 47(4):597–611, 2003.

Andrew Kydd. When can mediators build trust? *American Political Science Review*, 100(3):449–462, 2006.

Thomas Princen. *Intermediaries in International Conflict*. Princeton: Princeton University Press, 1992.

Robert W. Rauchhaus. Asymmetric information, mediation, and conflict management. *World Politics*, 58(2):207–241, 2006.

Gary Sick. The partial negotiator: Algeria and the U.S. hostages in Iran. In Saadia Touval and I. William Zartman, editors, *International Mediation in Theory and Practice*, pages 21–66. Boulder: Westview, 1985.

Kjell Skjelsbæk and Gunnar Fermann. The UN Secretary-General and the mediation of international disputes. In Jacob Bercovitch, editor, *Resolving International Conflicts: The Theory and Practice of Mediation*, pages 75–104. Boulder: Lynne Rienner, 1996.

Alastair Smith and Allan Stam. Mediation and peacekeeping in a random walk model of civil and interstate war. *International Studies Review*, 5(4):115–135, 2003.

Lesley G. Terris and Zeev Maoz. Rational mediation: A theory and a test. *Journal of Peace Research*, 42(5):563–583, 2005.

Saadia Touval and I. William Zartman. Mediation in international conflicts. In Kenneth Kressel and Dean C. Pruitt, editors, *Mediation Research: The Process and Effectiveness of Third-Party Intervention*, pages 115–137. San Francisco: Jossey-Bass, 1989.

Oran R. Young. Intermediaries: Additional thoughts on third parties. *Journal of Conflict Resolution*, 16(3):51–65, 1972.

I. William Zartman and Maureen R. Berman. *The Practical Negotiator*. New Haven: Yale University Press, 1982.