Bargaining and Learning While Fighting

Robert Powell University of California

Much of the existing formal work on war models the decision to go to war as a game-ending, costly lottery. This article relaxes this assumption by treating war as a costly process during which the states run the risk of military collapse. The model also allows for uncertainty over either the cost of fighting or the distribution of power. The analysis makes four contributions to the growing costly-process literature: (i) the present model provides a more general treatment of the learning process that occurs when states are uncertain about the distribution of power, (ii) it explicitly compares the bargaining and learning processes for the two different sources of uncertainty, (iii) it suggests a way to empirically distinguish wars arising from these two sources, and (iv) it shows that the equilibrium dynamics of informational accounts of war may be quite sensitive to the underlying bargaining environment through which information is conveyed.

Recent work on the causes of war raises two related issues. First, both formal and nonformal analyses see uncertainty as a fundamental cause of war. But ordinary-language analyses typically emphasize uncertainty about the distribution of power.[1] Blainey, for example, argues that "wars usually begin when fighting nations *disagree* on their relative strength" (1973, 122). By contrast, many formal analyses focus on uncertainty over the costs of fighting.[2] This emphasis largely reflects tractability considerations in that introducing uncertainty via costs rather than the distribution of power has led to models that have proved much easier to analyze.

These different sources of uncertainty pose important theoretical and empirical questions: Do these different sources of uncertainty lead to fundamentally different dynamics? And what empirical patterns, if any, distinguish wars arising primarily from one source of uncertainty from those that are largely due to the other?

A second, related issue is that until recently most formal studies of the causes of war treated the decision to go to war as a game-ending move (e.g., Fearon 1995; Powell 1999; Schultz 1999; Smith 1998b). Strategic interaction stops once the states decide to fight, and each state typically receives a payoff that reflects the distribution of power. In the language of bargaining theory, this work models war as an "outside option" which, if exercised, ends the bargaining.

This approach has proved to be quite fruitful. Informational accounts based on costly-lottery models offer a more coherent theoretical explanation of the origins of war than do many other explanations based on anarchy or preventive war (Fearon 1995) or on the offense-defense balance or relative-gains concerns (Powell 2002, 11–13). These models also cast considerable doubt on traditional ideas about the tendency of states to balance (Powell 1999; Werner 2000). Moreover, analyses that model war as a game-ending move have demonstrated the theoretical and empirical importance of taking selection effects into account in, for example, understanding the relation between deterrence and *ex ante* indicators like alliances. Fearon (1994b), for instance, provides a natural explanation of the negative correlation between the existence of an alliance and deterring an attack. And, costly-lottery formulations may help explain why it has been very difficult to find any stable empirical relation between the distribution of power and the likelihood of war (Powell 1996; Wagner 1994).

However, the costly-lottery simplification comes at a cost. It necessarily limits the scope of any analysis to the

Robert Powell is Robson Professor of Political Science, University of California, Berkeley, Berkeley, CA 94720-1950 (rpowell@socrates.berkeley.edu).

I am grateful to James Morrow, Leo Simon, Alastair Smith, R. Harrison Wagner, and to the participants at the conference on the "Political Economy of Conflict," Yale University, March 23–24, 2001, for many helpful comments and criticisms. This work was assisted by a grant from the National Science Foundation (SES-9911075).

[1]Goemans (2000) surveys this work.

[2]See, for instance, Fearon (1994a, 1995), Kydd (1999, 2001), and Powell (1996a, 1996b, 1999). Morrow (1989) provides an early effort to model uncertainty over capabilities. Powell (2002) surveys bargaining models of war.

American Journal of Political Science, Vol. 48, No. 2, April 2004, Pp. 344–361

 ISSN 0092-5853

origins of war. These models cannot be used to study the dynamics of intrawar conflict and bargaining, e.g., the relation between the states' negotiating positions and battlefield outcomes (Goemans 2000). This limitation by itself is a good reason to try to relax the costly-lottery assumption. But Wagner (2000) argues that there is an even more compelling reason. He charges that this assumption "can only lead to misleading conclusions" even about the origins of war (2000, 469).

The desire to relax the game-ending, costly-lottery assumption—reinforced by Wagner's charge—has spawned what might be called a second wave of formal work on war (Filson and Werner 2002a,b; Kim 2002; Slantchev 2003, 2002b; Smith 1998a; Smith and Stam 2003; and Wagner 2000). This second wave treats war as a costly process during which states can continue to bargain while they fight. War, in other words, is being modeled as an "inside option."[3]

The present study sees war as a bargaining process during which the states run a risk of military collapse. (As elaborated below, this is closest in spirit to Wagner's (2000) formulation.) The bargaining in the model begins before there is any fighting, continues during the fighting, and concludes only when the states agree to a settlement or one of them collapses. The model developed here also allows for uncertainty over either the cost of fighting or about the distribution of power.

The analysis makes four contributions. First, it offers a more general treatment of the learning process that occurs when there is uncertainty over the distribution of power. Unlike Smith and Stam (2003), the present analysis incorporates the strategic transmission of information into the learning process and does so in a much less restrictive setting than Filson and Werner (2002a). (The latter assume that a state is unsure whether its adversary is either strong or weak, i.e., there are only two types, and that there can be no more than two rounds of bargaining.)

Second, allowing for uncertainty either over the cost of fighting or about the distribution of power makes it possible to explicitly compare the bargaining and learning processes in these two cases. This comparison shows the bargaining dynamics are quite similar. Regardless of the source of uncertainty, the satisfied state makes a series of ever more attractive concessions that "screen" the dissatisfied state according to its type. Higher-cost types, i.e., those for whom fighting is more costly, settle sooner on worse terms than lower-cost types when there is uncertainty about the cost of fighting. In a like manner, weaker or less powerful types settle sooner than stronger types if there is uncertainty over the distribution of power. The comparison also shows that Wagner's (2000) conjectures about the differences between the learning processes in these two cases are only partially correct.

That the screening dynamics are similar regardless of the source of uncertainty poses an empirical challenge. Given the prevalence of both types of uncertainty in international politics, how can one empirically distinguish wars resulting primarily from one source of uncertainty or the other? The third contribution is to suggest a way. States in many situations can make offers more quickly than they can prepare for and fight battles. This difference affects the screening processes. Screening can take place without actually fighting any battles if the states are uncertain about each other's costs and can make offers very rapidly. If, by contrast, the states are uncertain about the distribution of power, then screening requires that the states actually fight battles regardless of how fast they can make offers. This suggests that crises arising out of uncertainty over costs or resolve are likely to be settled more quickly and short of large scale fighting than are crises arising out of uncertainty over the distribution of power.

One of the reasons for trying to relax the costly-lottery assumption is that costly-process models can be used to study the factors that affect the duration of war. The fourth contribution is to underscore the potential influence of the bargaining environment (e.g., how fast states can make offers) in models where fighting results from informational problems. Indeed, the effects of the bargaining environment in the present model swamp other factors (at least in the limit as the time between offers goes to zero) that might be thought to affect duration such as the cost of fighting or the distribution of power. (See, for example, Bennett and Stam 1996). This suggests that future efforts to use informational asymmetries to explain war and, especially, long wars have to pay a great deal of attention to the mechanism through which the actors convey information (e.g., the bargaining environment or protocol). In particular, these accounts have to explain why it takes so long to resolve the informational asymmetry.

The next two sections specify the costly-process model and describe its relation to other formulations. The third section discusses the bargaining and learning that takes place in equilibrium when there is uncertainty either over the cost of fighting or the distribution of power. The fourth section extends the basic model by allowing multiple offers between battles. The appendix contains proofs of key results.

[3]In contrast to an outside option, an "inside" option keeps the bargaining going for at least one more period and defines what payoffs the bargainers get during that period. See Muthoo (1999) for a discussion of inside and outside options.

FIGURE 1 Models of War

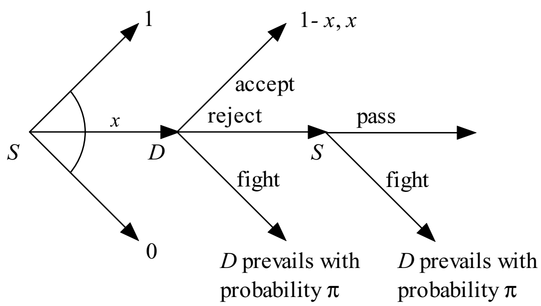

(a) A round of the costly-lottery game.

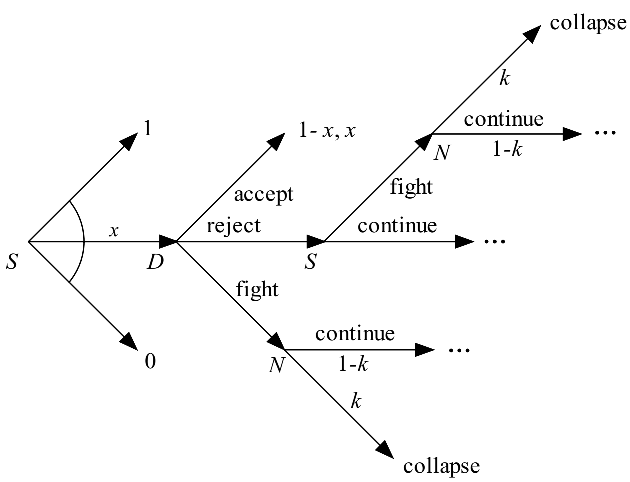

(b) A round of the costly-process game.

A Costly-Process Model of Conflict

This section presents a costly-process model which generalizes the costly-lottery approach. In a typical costly-lottery game, two states, say S and D, bargain about revising the territorial status quo, q.[4] Figure 1(a) depicts a round of the game. S begins the game by making an offer, $x \in [0, 1]$, to D who can accept the offer, reject it, or go to war to change the territorial status quo. Accepting ends the game with the territory divided as agreed. If D fights, the game ends and one state or the other is eliminated. More precisely, the game ends in a lottery in which both states pay a cost of fighting; D obtains all of the territory and eliminates S with probability π; and S eliminates D and thereby gains all of the territory with probability $1 - \pi$. The parameter π represents the distribution of power. If, for example, $\pi = \frac{1}{2}$, then each state has the same chance of winning and there is an even distribution of power. If, alternatively, π is close to one or zero, D or S has a preponderance of power. One can think of the game-ending lottery as a fight to the finish in which either S collapses and is unable to offer further resistance with probability π or D collapses and is unable to offer further resistance with probability $1 - \pi$. If D rejects S's offer, then S can attack or pass. Attacking again ends the game in a lottery. Passing ends the round, and another begins with S making an offer.[5]

In the costly-process game, fighting only creates a risk of collapse. If neither state collapses, the bargaining and fighting can continue until the states reach agreement or until one of them eventually does collapse. Assume, more specifically, that if S and D fight in a given round, then this fighting generates a risk k_S that S collapses in that round and a risk k_D that D collapses. These risks are assumed to be positive but possibly very small. Fighting is also costly, and S and D pay costs s and d each time they fight.[6]

Figure 1(b) illustrates a round of the game. As before, state S begins by proposing a territorial division $x \in [0, 1]$. D can accept this offer, reject it, or fight. If D accepts, the game ends with D and S in control of x and $1 - x$, respectively. If D fights, the states pay costs s and d, respectively, and the game ends with probability $k \equiv 1 - (1 - k_D)(1 - k_S)$, which is the probability that one or both of the states collapse. If only S collapses, D obtains all of the territory. If D alone collapses, S gets everything. If both collapse, the territorial status quo q remains in place. Finally, if neither state collapses, the current round ends, and the next begins. If D rejects S's offer without fighting, then S has the option of fighting or not. Not fighting ends the round and the next begins. If S fights, one or both states collapse with probability k and the game ends. If neither state collapses, the round ends and the next starts with a new offer from S.

It will be useful to define the distribution of power in the costly-process game. The distribution of power π in the costly-lottery formulation is the probability that D prevails in a fight to the finish (which is the only kind of

[4]See Fearon (1995) and Powell (1999) for extensive discussions of the costly-lottery model.

[5]There are several variants of the costly-lottery model depending on which bargainers can make offers and how many they can make. At one extreme, Fearon (1995) allows only one bargainer to make a single take-it-or-leave-it offer. At the other extreme, Powell (1999) allows both bargainers to make as many offers as they wish. In the costly-process model discussed here, only one bargainer can make offers but that bargainer can make as many offers as it pleases. One-sided, infinite-offer models are more general than take-it-or-leave-it models and easier to analyze than alternating-offer models.

[6]In the present analysis, s and d are taken to be the tangible costs of war measured in lives and economic losses. These costs will be discussed more fully below. Alternatively, one might think of s and d as more subjective measures of the states' resesolve or willingness to fight. The larger s and d, the lower the states' utilities to fighting, and the less resolute or willing to fight they are.

fight that can happen in that model). But there is another equivalent way of interpreting π: It is the expected territorial division resulting from a fight to the finish. That is, π also equals *D*'s expected share of the territory, i.e., $\pi = (1)\pi + 0(1-\pi)$ where, recall, *D* obtains all of the territory with probability π and loses everything with probability $1-\pi$. Paralleling this, define the *distribution of power p in the costly-process model* to be *D*'s expected territorial division resulting from a fight to the finish. This is:

$$p \equiv \frac{k_S(1-k_D)}{k} + \frac{qk_Sk_D}{k}. \quad (1)$$

Note that the distribution of power shifts in favor of *D* as its probability of collapse decreases, i.e., p is strictly decreasing in k_D.

The payoffs in the game are defined formally in the appendix. Suppose, less formally, that the states agree at time a to territorial division x after fighting at times $f_1, f_2, \ldots, f_n$. The expected payoff to this outcome is the payoff if the game ends when the states fight at f_1 weighed by the probability that the game ends at that time, plus the payoff if the game ends at time f_2 weighted by the probability that the game ends at that time, and so on through the n-th time that the states fight, plus the payoff to agreeing to x at time a weighted by the probability that the states did not collapse at times $f_1, \ldots, f_n$.

Finally, the model assumes one-sided incomplete information. *S* is uncertain either about *D*'s cost of fighting, d, or about *D*'s probability of collapse, k_D. Because p is strictly decreasing in k_D, uncertainty over k_D is equivalent to uncertainty over the distribution of power p (and working with uncertainty over the former is much easier). Although *S* is uncertain of *D*'s cost or of the distribution of power, *S* believes that *D* is dissatisfied, i.e., *D* prefers fighting to the finish to living with the status quo.[7]

Other Costly-Process Models

Recent work has relaxed the costly-lottery assumption in several different ways that highlight complimentary aspects of the underlying process and reflect different modeling trade offs.[8] The present formulation is closest to Wagner (2000). He also bases his analysis on a model in which rejecting an offer generates a risk that the game will end in that round because one state or the other collapses. Wagner, however, does not formally derive the equilibria of the game when there is asymmetric information, and, as will be discussed more fully below, some of his conjectures about the equilibrium dynamics are only partially correct.

Smith and Stam (2003) model the costly process differently, emphasizing the way that the probability of ultimate success in a fight to the finish can shift back and forth with the tide of battle. Building on Smith (1998a), they assume that two states, say *S* and *D*, are bargaining and that there are n "forts" which are initially divided between the states. *S* begins the game by making an offer to *D* who can accept or reject. Accepting ends the game on the agreed terms. If *D* rejects, the states fight with *D* winning a fort with probability η and losing a fort with probability $1-\eta$. The game continues in this way with *S* making all of the offers until the states reach agreement or until a state collapses because it no longer controls any forts.

The states also disagree about the probability η. This disagreement is not due to differences in the states' information about, say, their respective military capabilities but to a more basic subjective difference about the balance of power. Formally, Smith and Stam assume that the states have complete information but heterogeneous priors. This formulation contrasts with the informational approach generally taken in economics and in the existing formal work on war. (Smith and Stam is the only exception.) In this approach, actors are assumed to share a common prior and any differences in their beliefs are attributed to informational differences.[9]

Slantchev (2003) also uses Smith's (1998a) basic model of warfare but allows for alternating offers in the context of one-sided uncertainty about the distribution of power η. Following the informational approach, he attributes the differences in the states' beliefs about the distribution of power to informational differences. In this case, one has state has private information about its military capability.

Filson and Werner (2002a,b) underscore the resource constraint states face when fighting. In their formulation, the game begins with each state in control of a limited amount of resources. As in Smith and Stam and Slantchev, the states fight a battle whenever an offer is rejected with *D* wining with probability ψ and losing with probability $1-\psi$. Fighting a battle also consumes resources (although the amounts may differ depending on whether a state wins or loses), and a state collapses when it runs

[7] *S*'s beliefs are formalized in the appendix. The assumption that *D* is dissatisfied eliminates a mulitplicity of equilibria.

[8] See Wittman (1979) for an earlier, path-breaking effort to model conflict.

[9] An extended discussion of the common priors assumption is beyond the scope of the present analysis. (Morris 1995 provides a good overview.) Suffice it to say here that this assumption is extremely strong and at best justified on methodological grounds: among other things, it helps to discipline arguments. But it is also well worth trying to relax this assumption in a disciplined way.

out of resources. Thus, the bargaining continues until the states reach agreement or until one of the states runs out of resources. There is also one-sided uncertainty about ψ.

These formalizations of fighting embody different modeling trade offs. The costly process in Filson and Werner (2002a) is the only one that explicitly models the idea that resources are limited and consumed by fighting. However, their treatment of bargaining is the simplest and most restrictive: D can be one of only two possible types (i.e., S is unsure whether D's probability of prevailing in a battle, ψ, is high or low), and the states' resources are so constrained that there can be no more than two battles.[10] Smith and Stam (2003) and Slantchev (2003) trade a somewhat simpler formulation of the costly process for a richer treatment of the dynamics of bargaining. They abstract away from the resource constraint but do not impose exogenous limits on the number of offers. Slantchev's alternating-offer formulation also allows him to study signaling, albeit in a limited setting where D is one of only three possible types.[11] The costly process in the present formulation is the simplest, abstracting away from both the effects of the ebb and flow of battle and from any resource constraints. But the analysis places fewer restrictions on S's uncertainty about D, and this allows for a more general treatment of the learning and updating that occurs while fighting.[12] And, as will be seen below, the model can be readily extended to allow for multiple offers between battles.

Bargaining and Learning in Equilibrium

A satisfied state in a costly-lottery game faces a trade off when deciding what to offer the dissatisfied state. The more the satisfied state concedes, the better the chances of a settlement but the worse the terms of that settlement. The optimal offer that resolves this trade off typically entails some risk that the dissatisfied state will reject the offer and go to war.

In the costly-process game, the satisfied state confronts the same trade off when making its first offer, and the optimal offer usually entails some risk of rejection and fighting. But this fighting generally does not end the game. With probability $1 - k$, neither state collapses during the first round of fighting, and the bargaining continues. The satisfied state updates its beliefs based on the facts that (i) the dissatisfied state rejected the opening offer and that (ii) the dissatisfied state fought one round without collapsing. In light of these updated beliefs, the satisfied state makes a second offer which, like the first, usually carries some risk of rejection. The bargaining continues in this way until the dissatisfied state accepts an offer; one of the states collapses; or there is so little uncertainty left that S's optimal offer is to concede enough to buy D off and thereby end the bargaining.

These offers define the basic screening process that characterizes the equilibrium, and this process is the same regardless of the source of uncertainty. When there is uncertainty over D's cost of fighting, high-cost types accept smaller offers and settle earlier than do low-cost types who hang on longer and run greater risks of collapse in order to obtain larger concessions. More specifically, S's offers $x_1 < x_2 < x_3 < \cdots < x_n$ induce a set of cut points $\bar{d} > \tilde{d}_1 > \tilde{d}_2 > \cdots > \tilde{d}_n = \underline{d}$ such that the dissatisfied state counters all of the satisfied state's offers by fighting until it accepts x_j (assuming that neither of the states has already collapsed) if its actual cost of fighting d is between $\tilde{d}_{j-1}$ and $\tilde{d}_j$. Similarly, S's concessions when there is uncertainty over the distribution of power induces a set of cut points $\bar{k}_D > \tilde{k}_1 > \tilde{k}_2 > \cdots > \tilde{k}_m = \underline{k}_D$ such that the dissatisfied state counters the satisfied state's offers by fighting until accepting x_j if the dissatisfied state's probability of collapse k_D is between $\tilde{k}_{j-1} > \tilde{k}_j$. Stating the results more formally:

Proposition 1. *Suppose that S is satisfied, D is dissatisfied, and that S's beliefs about D's cost of fighting satisfy the technical assumptions specified in the appendix. Then there is a generically unique perfect Bayesian equilibrium. In it, S makes a finite number of offers which are strictly increasing. D responds to these concessions by fighting until ultimately accepting an offer. The lower D's cost d, the longer D fights and the better terms it obtains, i.e., d fights at least as long as* d' *if* $d < d'$.

Proposition 2. *Suppose that S is satisfied, D is dissatisfied, and that S's beliefs about D's probability of collapse satisfy the technical assumptions specified in the appendix. Then there is a generically unique perfect Bayesian equilibrium. In it, S makes finitely many offers which are strictly increasing. D responds to these concessions by fighting until*

[10] Filson and Werner (2002b) subsequently relax these constraints to some extent in the context of specific numerical examples.

[11] Signaling occurs when a state with private information makes offers. All of the other asymmetric-information models discussed here focus on screening, i.e., the state with private information receives offers but cannot make them.

[12] Formally, D's type is not restricted to a finite number of potential types. Rather, D's type is assumed to be distributed according to a continuous distribution which satisfies a few technical restrictions (which are described in the appendix).

ultimately accepting an offer. The more powerful D, the longer it fights and the better terms it obtains, i.e., k_D fights at least as long as k'_D if $k_D < k'_D$.

The demonstration of these results is extremely tedious. The appendix outlines the basic logic of the overall argument and proves key lemmas. Powell (2003) provides a more complete analysis.

Although the screening processes are broadly similar regardless of the source of uncertainty, there are subtle differences in the learning processes underlying them. Wagner (2000) argues that the bargaining and learning that take place when there is uncertainty over the distribution of power are different and simpler than they are in standard economic models of bargaining in which, say, a buyer has private information about how much he is willing to pay for an object for sale.

> In the standard bargaining models used in economics, private information explains why agreement is not immediate, but the only way that bargainers have to reveal their private information is by temporarily refusing to agree. Since they determine whether and how long they will hold out, they have an incentive to use this decision to misrepresent their private information. As a result, the signals they give by deciding whether to hold out or not are noisy and can be interpreted only by taking into account the strategic incentives of the bargainers . . .
>
> Bargaining in the context of war is different, in that fighting is a source of information that is much less subject to manipulation by adversaries (2000, 478).

To the extent that Wagner's argument holds, it goes beyond the economic models of bargaining. It also applies to costly-process models of war in which states have private information over the cost of fighting as in the model above. (See Kim 2002 for another example.) The only way that states can signal their costs in these models, as in the standard bargaining models in economics, is by refusing to agree.

However, the analysis underlying Proposition 2 shows that Wagner is only partially correct. When there is asymmetric information about the distribution of power, fighting does convey some information that, in a certain sense, is less subject to strategic manipulation. But this less manipulable channel does not replace the more manipulable channel present in the standard models. Both channels are present.

In asymmetric-information models, one actor learns about another actor or, more specifically, about another actor's type by watching how that actor behaves and then updating its beliefs about that actor's type in light of the observed behavior. Suppose, for example, that a seller makes an offer that one type of buyer, say b, is more likely to reject than type b' is. If the buyer subsequently does reject this offer, then the seller reasons backwards from this behavior and, applying Bayes' law, becomes more confident that it is facing type b who was more likely to reject the offer.

In what Wagner calls the standard model, learning depends on different types behaving differently. If two types behave identically, their behavior does not reveal anything that helps another actor distinguish between them. Suppose, for instance, that b and b' are equally likely to reject an offer that is ultimately rejected. The seller learns nothing about the relative likelihood of these two types from this rejection because both types would have behaved in the same way.

In contrast to the standard models, fighting when there is uncertainty over the distribution of power does convey information about types even if they behave in the same way. Suppose, for example, that the satisfied state is trying to distinguish between types k_D and k'_D where k_D is less likely to collapse than k'_D ($k_D < k'_D$) and therefore is more powerful. At the outset of the game, S believes that D's probability of collapse is distributed according to Θ with density θ. Consequently, the odds that S is facing k_D rather than k'_D are given by the ratio $\theta(k_D)/\theta(k'_D)$. Assume further that in equilibrium both of these types reject S's initial offer. In a standard model, the odds of facing k_D or k'_D would remain $\theta(k_D)/\theta(k'_D)$ because both types behave in the same way. But the fact that both types fought *and did not collapse* reveals information about them even though they behaved identically.

To see that this is the case, observe that the chances that D's probability of collapse is k_D given that it fought once and did not collapse is the probability of not collapsing, $(1 - k_D)$, times the prior $\theta(k_D)$. Consequently, the odds of k_D versus k'_D given that D did not collapse in the first round of fighting are $(1 - k_D)\theta(k_D)/[(1 - k'_D)\theta(k'_D)]$. These updated odds are greater than the initial odds $\theta(k_D)/\theta(k'_D)$ because $k_D < k'_D$. Thus, the satisfied state becomes more confident that it is facing the more powerful type k_D relative to k'_D even though both types behaved identically. Indeed, the odds of facing k_D increase with each round of fighting. If these types have fought r times without collapsing, the odds rise to $(1 - k_D)^r\theta(k_D)/[(1 - k'_D)^r\theta(k'_D)]$. This is the sense in which fighting conveys information that is less subject to strategic manipulation.

Fighting, however, is only one of the ways that information is conveyed when there is asymmetric information. Dissatisfied types also have an incentive to misrepresent their types by delaying an agreement in order to appear "tougher" just as they do in standard models. To verify this, consider S's beliefs about D when making the $r+1$-st offer, x_{r+1}. The proof of Proposition 2 shows that these beliefs depend on the prior $\Theta(k_D)$ updated in light of the facts that D fought r times without collapsing and that k_D must be no more than k_r (otherwise D would have accepted a previous offer).

To formalize this updating, define the posterior distribution $\Theta^r(k_D)$ to be the prior distribution Θ given that D has fought r times without collapsing. This updated distribution is $\Theta^r(k_D) = [\int_{\underline{k}_D}^{k_D}(1-z)^r\theta(z)\,dz]/[\int_{\underline{k}_D}^{\bar{k}_D}(1-z)^r\theta(z)\,dz]$. Then, S's beliefs when making the $r+1$-st offer are given by Θ^r conditional on $k_D < k_r$, i.e., by $\Theta^r(k_D)/\Theta^r(k_r)$.

The fact that S's beliefs depend on both the distribution Θ^r and the cut point k_r means that information is conveyed through two channels when there is uncertainty over the distribution of power. Θ^r reflects the first channel. It depends solely on the prior Θ and the number of previous rounds of fighting. The cut point k_r captures the effects of the second channel. This cut point is the type that is indifferent between accepting the r-th offer and fighting one round before accepting the next offer. But, S's next offer depends on its beliefs when making that offer, and these beliefs depend on the cut point k_r. Hence, which types hold on for the next offer and what S offers next are strategically interdependent. This means that the signals dissatisfied types send "by deciding whether to hold out or not are noisy and can be interpreted only by taking into account the strategic incentives of the bargainers" (Wagner 2000, 478) even when there is uncertainty about the distribution of power.

In sum, bargaining while fighting when there is uncertainty or, more precisely, asymmetric information about the distribution of power does not substitute one channel of conveying information which is less manipulable for another. This type of bargaining combines both channels.

Multiple Offers Per Battle and the Importance of the Bargaining Environment

The preceding suggests that it may be quite difficult empirically to distinguish conflicts that arise because of uncertainty over costs or from those due to uncertainty over the distribution of power. Propositions 1 and 2 show that the basic screening processes are the same in the model regardless of the source of uncertainty and, consequently, cannot be used to differentiate between these sources.[13] Although the learning processes are subtly different, it seems doubtful that this difference could provide a basis for distinguishing between these sources given the limitations of the data that generally exist in international relations. This section proposes one way of differentiating between these two sources of uncertainty based on the idea that states can often make offers much more quickly than they can prepare for battle.[14] The analysis also demonstrates that the bargaining environment in informational explanations of war may play a critical role in shaping the dynamics of conflict, e.g., the duration of the fighting. The potential significance of these effects suggests that empirical studies need to take these effects into account.

To modify the model to allow for multiple offers between battles, suppose that S can make m offers between battles. That is, the states pay "mobilization" or "preparation" costs if the dissatisfied state "fights" in Figure 1(b) in response to the first $m-1$ offers by preparing to fight. (The substantive interpretation of these costs and their relation to the cost of fighting in the one-offer-per-battle game are discussed below.) But the states do not actually fight and therefore run no risk of collapse. If, however, the dissatisfied state fights in response to the m-th offer, the states do run a risk of collapse.

The ability to make multiple offers before having to run the risk of collapse affects the screening processes differently. If there is uncertainty over costs and if the number of offers between battles, m, is sufficiently large, then the states settle before they ever fight. If, by contrast, the states are uncertain about the distribution of power, screening requires that there must be some chance that the states fight regardless of the number of offers per battle. This suggests that crises driven by uncertainty over costs will frequently settle short of large scale fighting while those driven by uncertainty over the distribution of power are more likely to entail significant fighting.

More formally, suppose there are m offers per battle where the time between offers is $\Delta \equiv 1/m$ (and, therefore, the total time between battles is $m\Delta = 1$ as it is in the one-offer-per-battle game above). S begins the game by making an offer which D can accept or counter by fighting

[13]As might be expected, the comparative statics also appear to be quite similar although I have not been able to characterize all of them explicitly.

[14]This, of course, depends on the existing military and communications technologies. When proposals had to be carried by ship, bargaining may not have been faster. The Treaty of Ghent ending the War of 1812 was signed on December 24, 1814; American envoys carrying the treaty left London on January 2, 1815; and the treaty arrived in Washington, D.C., for ratification on February 14. Meanwhile, British and Amerian forces fought the Battle of New Orleans on January 8 (Hickey 1989).

or waiting. If D waits, S decides whether to fight or wait after which the round ends. If either state decides to fight, S and D pay costs s_m and d_m, but there is no risk of collapse. The game continues in this way through $m - 1$ offers. If, however, the states fight in the m-th round, then S and D collapse with probabilities k_S and k_D, respectively. If neither collapses, there is no further risk of collapse until S makes the $2m$-th offer. The game continues this way until the states agree on a settlement or one of them collapses. Finally, the discount factor and the states' payoffs are adjusted in order to make the games comparable regardless of the number of offers between battles.[15]

Now consider the equilibrium screening process in the m-offer game in which S is uncertain of D's cost. Because D pays this cost whether or not the states generate any risk of collapse, S, not surprisingly, can still use its offers to screen the dissatisfied state even when rejecting those offers entails no risk of collapse. Furthermore, Proposition 3 shows that the bargaining satisfies the Coase conjecture which posits that the outcome of the bargaining will be efficient, i.e., the bargaining is virtually certain to end almost immediately if the time between offers is small enough (or, equivalently, if the number of offers m is large enough).[16]

Proposition 3. *For any* $\varepsilon > 0$, *there exists an M such that the probability that the bargaining ends by time* ε *is greater than* $1 - \varepsilon$ *whenever the number of offers-per-battle is at least M*.

Proof: See the appendix for a sketch.

[15]Since the time between offers is shorter in the m-offer game, the states discount less between offers and the discount factor δ_m in the m-offer game is $\delta_m \equiv \delta^{1/m}$. The states' per-offer costs, s_m and d_m, must also be amended so that the total cost of preparing for and fighting one battle remains s and d. That is, the total cost of paying s_m for m-rounds, $\sum_{j=0}^{m-1} \delta_m^j s_m$, must equal s. This implies $s_m = s(1 - \delta_m)/(1 - \delta)$ and similarly for d_m. Finally, a state's per-period payoff, x_m, to controlling a fraction x of the territory in the m-offer game must be normalized so that it is the same as the payoff to controlling x for one period in the one-offer game. This leaves $x_m = x(1 - \delta_m)/(1 - \delta)$.

[16]The Coase conjecture, stationary strategies, and the fact that S never makes offers that are sure to be rejected (as described in Propositions 1 and 2) are closely related. Suppose D's strategy is stationary in the sense that D's response to an offer of x is independent of the previous history of offers leading up to x as long as x is larger than any previous offer. Gul and Sonnenschein (1988) then show (in a buyer-seller game) that stationarity implies that the Coase conjecture holds. It is also clear that S will never make an offer which is sure to be rejected. For S could always do better by skipping that offer, because stationarity ensures that skipping an offer does not affect D's response to subsequent offers. See Fudenberg and Tirole (1991, 401) and Gul, Sonnenschein, and Wilson (1986) for further discussion of the Coase conjecture.

The screening dynamics are different if S is uncertain of the distribution of power. If there is any bargaining at all, then there must be some chance that the states fight at least one battle. If, that is, S does not initially offer enough to buy off the toughest type $\underline{k}_D$ (and, therefore, any other type), then there must be a positive probability that the states actually fight a battle. (The appendix also establishes a condition sufficient to ensure that S does not immediately end the game by offering enough to buy $\underline{k}_D$ off at the outset of the game.)

Proposition 4. *Suppose S is uncertain of the distribution of power in the m-offer game. Then the game either ends immediately with S's offering enough to buy off the toughest type it might be facing, namely* $\underline{k}_D$, *or there is a positive probability that the states fight at least one battle.*

See the appendix for a formal discussion of this result. To sketch the basic argument, suppose S's first offer is x_1 and $\underline{k}_D$ strictly prefers waiting for a subsequent offer, say x_j, that S makes no later than round m, i.e., before any fighting takes place. This leads to a contradiction. Because all types k_D pay the same cost d_m and do not have to run any risk of collapse in order to obtain x_j, the payoff to accepting x_1 or waiting for x_j is the same for all types. Hence, all k_D are sure to reject S's initial offer x_1 and wait for x_j because $\underline{k}_D$ strictly prefers waiting for x_j. But this is a contradiction as it is straightforward to show that S never makes an offer in equilibrium which is sure to be rejected because S could do better by simply skipping that offer.[17]

Two conclusions follow from these results. First, to the extent that asymmetric information is the underlying cause, crises due to uncertainty over the distribution of power seem likely to last longer and entail more fighting than those growing out of uncertainty over costs.[18] Indeed, there may be little or no fighting in the latter case if states can make offers quickly.

The second, broader conclusion is that the bargaining environment through which actors convey information may play a critical role in informational accounts of war and, especially, of prolonged conflict. Costly-lottery models obscure this role, and its importance has heretofore not been fully appreciated in either theoretical or empirical work on war. For example, existing empirical studies of duration (e.g., Bennett and Stam 1996) tend

[17]If $\underline{k}_D$ is indifferent between x_j and x_1 and the game does not end immediately, then an atom of types must be holding out for x_j. But this too leads to a contradiction as S would be able to increase its payoff by offering slightly more than x_1 and settle immediately.

[18]For discussions about noninformational explanations of war, see Fearon (1995) and Powell (2002, 23–27).

to focus on factors like the states' total military capabilities and the distribution of power and do not consider the bargaining environment in which the underlying conflicts play out. But the results above show that features of the bargaining environment such as how fast states can make offers can swamp the effects that factors like the distribution of power have on the duration of war.

This second conclusion requires elaboration for it depends on a careful interpretation of "costs" and, more subtly, of the relation between formal results and their substantive meaning. In the one-offer-per-battle model, s and d conflate the economic costs of mobilizing and preparing for war with the costs of actually fighting, the loss of life and physical destruction. This is also true of other costly-process models (e.g., Filson and Werner 2002; Slantchev 2003; Smith and Stam 2003) and the costly-lottery models.

The multiple-offer model makes it possible, indeed requires, that these costs be separated. Let d_p be D's per-period cost of preparing to fight and d_f be its cost of fighting. Then the total cost of preparing for and fighting a battle, which is simply s in the one-offer game, can be broken down into the total cost of mobilizing in $m - 1$ rounds plus the cost of fighting in the m-th round. In symbols, $s = \sum_{j=0}^{m-2} \delta_m^j d_p + \delta_m^{m-1} d_f$ where $\delta_m = \delta^{1/m}$ is the comparable discount rate for the m-offer game (see footnote 15).

The discussion above assumed that d_p and d_f were equal, i.e., $d_p = d_f = d_m$. If, alternatively, these costs are unequal but directly related to each other with higher mobilization costs associated with higher costs of fighting, then S can still screen D during the mobilization phase (i.e., during the $m - 1$ offers preceding a battle) because continuing to mobilize in order to secure a better offer costs different types different amounts. Proposition 3, therefore, continues to hold.

Suppose, however, that d_p and d_f are completely unrelated and that the mobilization costs are the same for every type of D, i.e., what distinguishes different types of D are their costs of fighting d_f. If so, then S can no longer use d_p to screen D during the mobilization phase and results analogous to those in Proposition 4 obtain: There must be some chance of fighting if the game does not end immediately.[19]

All of this underscores the potential sensitivity of informational accounts of war to the bargaining environment—to the sources of uncertainty and the ability to resolve that uncertainty.[20] This sensitivity means that future informational efforts should either include some robustness checks for a particular formalization of the bargaining environment or they should ground that formalization empirically. For example, a satisfactory explanation of prolonged fighting that only holds if the states cannot make offers quickly must also explain in an empirically convincing way why the states cannot make offers quickly.

Conclusion

Modeling war as a game-ending move has proved enormously useful both theoretically and empirically. But this simplification necessarily limits the scope of the analysis to the origins of war. Investigating the strategic interaction inherent in the conduct of war requires more explicit models of the process of fighting.

The present study sees war as a process during which the states can continue to bargain while facing a risk of military collapse. The model also allows for uncertainty over the costs of fighting or over the distribution of power. In equilibrium, the satisfied state, regardless of the source of uncertainty, makes a series of ever more attractive offers. These offers screen the dissatisfied state with tougher types (i.e., stronger or low-cost types) fighting longer in order to secure a better agreement but at the cost of running a greater risk of military collapse. There are, however, subtle differences in the underlying learning processes associated with each source of uncertainty. When there is uncertainty over the distribution of power, fighting itself conveys information even about types that behave identically.

To distinguish between conflicts arising from these different sources, the analysis made an identifying assumption: States can make offers more quickly than they can prepare for and fight battles. In these circumstances, crises arising from uncertainty over costs are likely to be shorter and entail less fighting (and in the limit no fighting). By contrast, crises resulting from uncertainty over the distribution of power are likely to be longer and entail significant fighting.

[19]The results in Proposition 4 depend on the fact that k_D and k'_D obtain the same payoff to agreeing to x at time t_1 and the same payoff to continuing to mobilize in order to obtain y at some later time t_2 as long as there are no battles between t_1 and t_2. The same is true if each type of D has the same mobilization cost d_p but different mobilization cost than d_f.

[20]To highlight this sensitivity and complicate matters still further, the present model, like other costly-process models, follows the Rubinstein-Ståhl approach of in which the timing of the offers is imposed exogenously. An alternative approach endogenizes the timing of the offers. Admati and Perry (1987) develop this approach in a buyer-seller context, and Heifetz and Segev (2002) use this approach to model conflict.

Finally, the critical role that the bargaining environment plays in the present analysis and is likely to play in other informational accounts of war poses a challenge for future work. Future efforts to explain conflict, especially prolonged conflict, in terms of informational problems must give detailed attention to and a plausible empirical account of the mechanism through which the bargainers can and cannot convey information.

Appendix

This appendix formalizes some aspects of the game and outlines the proofs of Propositions 2, 3, and 4. (The basic logic establishing Proposition 1 is the same as that of Proposition 2.) The proofs of some key lemmas are also presented, but limitations of space prevent a more complete presentation. (See Powell 2003 for omitted proofs.)

Some Preliminaries

Some preliminaries are needed before sketching the arguments. Recall that the distribution of power, $p = k_S(1 - \tau)/k + qk_S\tau/k$, is D's expected territorial share resulting from a fight to the finish. (For expositional ease, it will be convenient to use the simpler notation τ instead of k_D to represent D's probability of collapse where $\tau \in [\underline{\tau}, \bar{\tau}]$.) The distribution of power can be interpreted in another way which will be useful in specifying the states' payoffs. Suppose the game is in a round, say m, and ends in that round because one or both of the states collapses in that round. Then, the distribution of power is also equivalent to the expected territorial division resulting from that collapse. To establish this, note that the probability that S alone collapses given that one of the states collapses is $k_S(1 - \tau)/k$, and the probability that S and D collapse given that one of them does is $k_S\tau/k$. Hence, the first term in the expression for p is the total territory (i.e., one) weighted by the conditional probability that S alone collapses, and the second term is D's status quo share q times the conditional probability that both states collapse. Thus, p is D's expected territorial share conditional on the game's ending in a specific round because one or both states collapses in that round.

To formalize the payoffs, suppose that the states agree at time a to territorial division x after fighting at times $f_1, f_2, \ldots, f_n$. The probability that the game ends when the states fight for the m-th time in period f_m is the probability that the states do not collapse during their first $m - 1$ fights and then do. This is $(1 - k)^{m-1}k$. D's payoff if the states collapse at this time is its return to having q through time f_m less paying d in rounds $f_1, f_2, \ldots, f_m$ when the states fight plus the payoff to having the expected territorial share p from time $f_m + 1$ on. This is $\sum_{j=0}^{f_m} \delta^j q - \sum_{j=1}^{m} \delta^{f_j} d + p \sum_{j=f_m+1}^{\infty} \delta^j$. The probability that the states do not collapse before agreeing to x at time a after fighting n times is $(1 - k)^n$, and D's payoff to this outcome is $\sum_{j=0}^{a-1} \delta^j q - \sum_{j=0}^{m} \delta^{f_j} d + \sum_{j=a}^{\infty} \delta^j x$. Pulling all of this together, D's expected payoff to agreeing to x at a after fighting at times $f_1, f_2, \ldots, f_n$ is: $[\sum_{m=1}^{n}(1 - k)^{m-1}k(\sum_{j=0}^{f_m} \delta^j q - \sum_{j=1}^{m} \delta^{f_j} d + p \sum_{j=f_m+1}^{\infty} \delta^j)] + (1 - k)^n[\sum_{j=0}^{a-1} \delta^j q - \sum_{j=0}^{m} \delta^{f_j} d + \sum_{j=a}^{\infty} \delta^j x]$. S's payoff is defined analogously.

A player-type τ of D is dissatisfied in the costly-process game if τ's payoff to fighting to the finish, $F_D(\tau)$, is greater than its payoff to living with the status quo. In symbols, τ is dissatisfied if $F_D(\tau) = \sum_{m=0}^{\infty}(1 - k(\tau))^j k(\tau)[\sum_{n=0}^{m}(q - d)\delta^n + p(\tau) \sum_{n=m+1}^{\infty} \delta^n] > q/(1 - \delta)$. This reduces to $p(\tau) - q > d(1 - \delta)/(\delta\, k(\tau))$.

Now let $\tilde{x}(\tau)$ denote τ's certainty equivalent of fighting to the finish. That is, $\tilde{x}(\tau)$ satisfies $\tilde{x}(\tau)/(1 - \delta) = F_D(\tau)$ which yields $\tilde{x}(\tau) = [\delta k(\tau) p(\tau) + (1 - \delta)(q - d)]/[1 - \delta(1 - k(\tau))]$. It is straightforward to show that D's certainty equivalent for fighting increases as its probability of collapse τ decreases, i.e., is $\tilde{x}(\tau)$ decreasing in τ. If, moreover, τ is dissatisfied, then τ's payoff to fighting to the finish is strictly higher than its payoff to living with the status quo: $\tilde{x}(\tau) > q$.

Formalizing S's beliefs, S believes that d is distributed over $[\underline{d}, \bar{d}]$ according to the cumulative distribution function $G(d)$ which has a bounded and continuous density $0 < \underline{g} \leq g(d) \leq \bar{g}$. That D is sure to be dissatisfied implies $\delta k(p - q) - (1 - q)d \leq 0$ for all $d \in [\underline{d}, \bar{d}]$. If S is uncertain about the distribution of power, then τ is distributed over $[\underline{\tau}, \bar{\tau}]$ according to $\Theta(\tau)$ which has a bounded density $0 < \underline{\theta} \leq \theta(\tau) \leq \bar{\theta}$. It will also be necessary to assume that the derivative of the density, $\theta'(\tau)$, is bounded at least in a neighborhood of $\underline{\tau}$. As before, Θ is common knowledge as is the fact that D is dissatisfied, i.e., $\delta k(\tau)(p(\tau) - q) - (1 - \delta)d \leq 0$ for all $\tau \in [\underline{\tau}, \bar{\tau}]$.

Finally, Lemma 3 below shows that S's beliefs along the equilibrium path are described by a cut point τ' and the number of previous fights r (along with the prior Θ). Call the pair (τ', r) the *state of the game*.

Basic Lemmas for Proposition 2

Proving Proposition 2 or the stronger version stated below takes three steps. The first establishes four lemmas that describe conditions that any perfect Bayesian equilibrium (PBE) of the game must satisfy. The second step uses these results to specify a dynamic programing problem and then shows that the solution to this problem is generically unique. Finally, a series of offers and cut points is a PBE-path if and only if it is a solution to the dynamic programing problem.

The first lemma shows that if $\hat{\tau}$ is the toughest type that S might be facing, then S never offers more than $\hat{\tau}$'s certainty equivalent of fighting to the finish, $\tilde{x}(\hat{\tau})$. The second lemma demonstrates that D never waits; it either accepts the current proposal or fights. Third, S's beliefs along the equilibrium path are truncations of Θ^r, i.e., the prior beliefs Θ conditioned on the number of previous fights r. Finally, there is a $\tau^* > \underline{\tau}$ such that S ends the game by buying off the toughest type by offering $\tilde{x}(\underline{\tau})$ whenever the state is bounded above by τ^*.

Lemma 1. *Fix a PBE and let $\hat{h}$ be any history, $\mu(\hat{h})$ be the support of S's beliefs at $\hat{h}$, $\hat{\tau} = \min\{\tau : \tau \in \mu(\hat{h})\}$ be the toughest type S might be facing at $\hat{h}$, and h be any history that passes through $\hat{h}$, i.e., $\hat{h} \subseteq h$. Then S never offers more than $\tilde{x}(\hat{\tau})$ at h.*

Lemma 2. *Consider any PBE and let h be any history (not necessarily on the equilibrium path) at which S makes an offer. Taking $\mu(h)$ to be the support of S's beliefs at h, the set of $\tau \in \mu(h)$ that wait in response to S's offer at h has measure zero.*

Lemma 3. *S's beliefs in period r along any PBE path are a truncation of $\Theta^r(\tau)$.*

Lemma 4. *There exists a $\tau^* > \underline{\tau}$ such that S offers $\tilde{x}(\underline{\tau})$ in any state (τ', r) if $\tau' \le \tau^*$.*

Defining a Dynamic Programming Problem

The next step in the proof of Proposition 2 is to specify a dynamic programing problem. Lemma 1 ensures that S never offers more than the certainty equivalent of the toughest type. Suppose, then, that S is exogenously constrained to make at most one offer before proposing $\tilde{x}(\underline{\tau})$. If the current state is (τ', r), S's payoff to offering x is:

$$V_{1|r}(x \mid \tau') \equiv \frac{1-x}{1-\delta}\left(\frac{\Theta^r(\tau') - \Theta^r(\tilde{\tau})}{\Theta^r(\tau')}\right) + \int_{\underline{\tau}}^{\tilde{\tau}} \left[1 - q - s + \delta k(\tau)\left(\frac{1-p(\tau)}{1-\delta}\right) + \delta(1-k(\tau))\left(\frac{1-\tilde{x}(\underline{\tau})}{1-\delta}\right)\right]\frac{d\Theta^r}{\Theta^r(\tau')}$$

where $\tilde{\tau}$ is the type that is indifferent between accepting x and waiting for $\tilde{x}(\underline{\tau})$. The first term is S's payoff if D accepts x weighted by the probability that D accepts. The second term is S's expected payoff to fighting one more period, namely, S's status quo share $1-q$, less the cost of fighting s, plus the discounted value of its territorial share if the game ends in collapse weighted by the probability of collapse $k(\tau)$, plus the discounted the payoff to settling on $\tilde{x}(\underline{\tau})$ weighted by the probability that the states do not collapse. (The integral reflects the fact that S does not know D's probability of collapse.) Now define $\sigma_{1|r}(\tau') \equiv \arg\max_{x\in[0,1]} V_{1|r}(x \mid \tau')$, $\underline{\sigma}_{1|r}(\tau') \equiv \min\{y : y \in \sigma_{1|r}(\tau')\}$, $\bar{\sigma}_{1|r}(\tau') \equiv \max\{y : y \in \sigma_{1|r}(\tau')\}$, $\hat{\sigma}_{1|r}(\tau')$ to be the convexification of $\sigma_{1|r}(\tau')$, i.e., $\hat{\sigma}_{1|r}(\tau') \equiv \{y : \underline{\sigma}_{n|r}(\tau') \le y \le \bar{\sigma}_{n|r}(\tau')\}$, and $V_{1\,|\,n}(-\tau') \equiv \max V_{1\,|\,n}(x|\tau')$.

Working backward, suppose that starting in state $(\tau', r-1)$, S can make no more than two offers before offering $\tilde{x}(\underline{\tau})$. Then S would maximize

$$V_{2|r-1}(x \mid \tau') \equiv \frac{1-x}{1-\delta}\left(\frac{\Theta^{r-1}(\tau') - \Theta^{r-1}(\tilde{\tau})}{\Theta^{r-1}(\tau')}\right) + \int_{\underline{\tau}}^{\tilde{\tau}} \left[1 - q - s + \delta k(\tau)\left(\frac{1-p(\tau)}{1-\delta}\right) + \delta(1-k(\tau))\,V^*_{1|r}(\tilde{\tau})\right]\frac{d\Theta^{r-1}}{\Theta^{r-1}(\tau')}$$

subject to the constraint that $\tilde{\tau}$ is indifferent between accepting x and fighting one round for some $y \in \hat{\sigma}_{1|r}(\tilde{\tau})$. (The game-theoretic interpretation of y is that it is S's expected next offer. Because S's next offer must maximize its payoff, i.e., be an element of $\sigma_{1|r}(\tilde{\tau})$, the expected next offer must be an element of $\hat{\sigma}_{1|r}(\tilde{\tau})$.) Continuing in this way, define $V_{n|r}(x|\tau')$, $V^*_{n|r}(\tau')$, $\sigma_{n|r}(\tau')$, $\underline{\sigma}_{n|r}(\tau')$, $\bar{\sigma}_{n|r}(\tau')$, and $\hat{\sigma}_{n|r}(\tau')$ recursively.

Now specify the correspondence $\gamma_{n|r}(\tau)$ to be τ's possible payoffs to fighting one round in order to obtain an element of $\hat{\sigma}_{n|r}(\tau)$:

$$\begin{aligned}\gamma_{n|r}(\tau) &\equiv \{(1-\delta)(q-d) + \delta k(\tau)\,p(\tau) \\ &\quad + \delta(1-k(\tau))y : y \in \hat{\sigma}_{n|r}(\tau)\} \\ &= (1-\delta)(q-d) + \delta k_S[1-\tau(1-q)] \\ &\quad + \delta(1-k_S)(1-\tau))\hat{\sigma}_{n|r}(\tau).\end{aligned}$$

Also, let $\underline{\gamma}_{n|r}(\tau) \equiv \min\{y : y \in \gamma_{n|r}(\tau)\} = (1-\delta)(q-d) + \delta k(\tau)\,p(\tau) + \delta(1-k(\tau))\underline{\sigma}_{n|r}(\tau)$ and $\bar{\gamma}_{n|r}(\tau) \equiv \max \gamma_{n|r}(\tau)$.

Finally, to ease the notational burden normalize $V_{n|r}(x \mid \tau')$ by letting $W_{n|r}(x \mid \tau') \equiv (1-\delta)\Phi^r(\tau')\times V_{n|r}(x \mid \tau')$ and $W^*_{n|r}(\tau') \equiv (1-\delta)\Phi^r(\tau')\,V^*_{n|r}(\tau')$

where $\Phi^r(\tau') \equiv \int_{\underline{\tau}}^{\tau'} (1-\tau)^r \theta(\tau)\, d\tau$ and, consequently, $\Theta^r(\tau') = \Phi^r(\tau')/\Phi^r(\bar{\tau})$. This leaves

$$\begin{aligned} &W_{n+1|r}(y|\tau') \\ &\quad = (1-y)[\Phi^r(\tau') - \Phi^r(\tilde{\tau})] \\ &\qquad + \int_{\underline{\tau}}^{\tilde{\tau}} \big[(1-\delta)(1-q-s) + \delta k(\tau)(1-p(\tau)) \\ &\qquad + \delta(1-\delta)(1-k(\tau))V^*_{n|r+1}(\tilde{\tau})\big](1-\tau)^r\theta(\tau)d\tau \\ &\quad = (1-y)[\Phi^r(\tau') - \Phi^r(\tilde{\tau})] + \big[(1-\delta)(1-q-s) \\ &\qquad + \delta(1-qk_S)\big]\Phi^r(\tilde{\tau}) - \delta(1-qk_S)\Phi^{r+1}(\tilde{\tau}) \\ &\qquad + \delta(1-k_S)W^*_{n-1|r+1}(\tilde{\tau}) \qquad\qquad \text{(A1)} \end{aligned}$$

where $\tilde{\tau}$ is the subsequent cut point and the second equation uses the facts that $1-k(\tau) = (1-k_S)(1-\tau)$ and $k(\tau)p(\tau) = k_S[1-\tau(1-q)]$.

Lemma 5 establishes useful properties about $W_{n|r}(y\,|\,\tau)$, $\hat{\sigma}_{n|r}(\tau)$, and $\gamma_{n|r}(\tau)$:

Lemma 5. *The following properties hold for all $r \geq 0$ and $n \geq 0$:*

(i) *$W_{n|r}(y\,|\,\tau)$ is continuous in y and τ and strictly increasing in τ.*

(ii) *The correspondence $\hat{\sigma}_{n|r}(\tau)$ is upper hemi-continuous and nonincreasing (i.e., $\underline{\sigma}_{n|r}(\tau) \geq \bar{\sigma}_{n|r}(\tau')$ for any $\tau' > \tau$), whereas $\gamma_{n|r}(\tau)$ is upper hemi-continuous and strictly decreasing (i.e., $\underline{\gamma}_{n|r}(\tau) > \bar{\gamma}_{n|r}(\tau')$ for any $\tau' > \tau$).*

(iii) *Consider S's offer of y in $W_{n+1|r-1}(y|\hat{\tau})$. Then, the next cut point $\tilde{\tau}_{n+1|r-1}(y)$ is a continuous, nonincreasing function of y defined by the unique τ for which $y \in \gamma_{n|r}(\tau)$. The expected next offer $e_{n|r}(y)$ is a continuous, nondecreasing function of y. This offer is the unique e such that $e \in \hat{\sigma}_{n|r}(\tilde{\tau}_{n+1|r-1}(y))$ and $\tilde{\tau}_{n+1|r-1}(y)$ is indifferent between accepting y and fighting one round for $e_{n|r}(y)$, i.e., $y = (1-\delta)(q-d) + \delta k(\tilde{\tau}_{n+1|r-1}(y))p(\tilde{\tau}_{n+1|r-1}(y)) + \delta(1-k(\tilde{\tau}_{n+1|r-1}(y)))e_{n|r}(y)$.*

Proof: Figure 4 illustrates the correspondences $\hat{\sigma}_{n+1|r-1}(\tau)$ and $\gamma_{n|r}(\tau)$ and the state $\tilde{\tau}_{n+1|r-1}(y)$ induced by offering $y \in \sigma_{n+1|r-1}(\tau')$ (where $\sigma_{n+1|r-1}(\tau')$ is assumed to be a singleton). Arguing by induction on n, conditions (i)–(iii) hold trivially for $n = 0$ and $r \geq 0$ as $W_{0|r}(y|\tau) \equiv (1-\delta)(1-\tilde{x}(\underline{\tau}))\Phi^r(\tau)$ and $\sigma_{0|r}(\tau) = \{\tilde{x}(\underline{\tau})\}$. Suppose they hold for n and any $r \geq 0$. Equation (A1) immediately implies that $W_{n+1|r}(y|\tau)$ is strictly increasing in τ. The continuity of W^*_n and $\tilde{\tau}_{n+1|r-1}$ also ensures that $W_{n+1|r}(y|\tau)$ is continuous in y. Thus, condition (i) holds for $n+1$.

FIGURE A.1 Solutions to $W_{n|r}(\tau)$

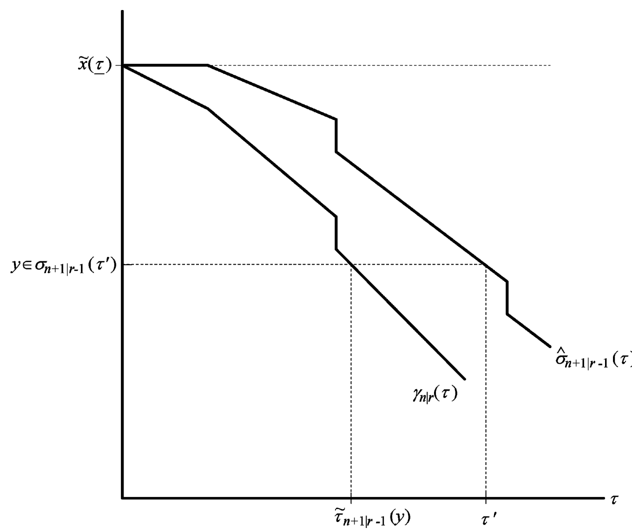

As for condition (ii), $W_{n+1|r}(y|\tau)$ is continuous and being maximized over the compact set $y \in [0,1]$ which varies continuously (or rather does not vary at all) with τ. Consequently, $\sigma_{n+1|r}(\tau)$ is well-defined and upper hemi-continuous. The convexification $\hat{\sigma}_{n+1|r}(\tau)$ is also upper hemi-continuous since $\sigma_{n+1|r}(\tau)$ is bounded for all τ and the graph of $\hat{\sigma}_{n+1|r}(\tau)$ is closed.

To see that $\hat{\sigma}_{n+1|r}(\tau)$ is nonincreasing, let $\tau' > \tau$. Since offering y induces the same cut point regardless of whether the current state is τ' or τ (see Equation A1), $W_{n+1|r}(y|\tau') - W_{n+1|r}(y|\tau) = (1-y)[\Phi^{r+1}(\tau') - \Phi^{r+1}(\tau)]$. Consequently,

$$\begin{aligned} &[1-\bar{\sigma}_{n+1|r}(\tau')][\Phi^{r+1}(\tau') - \Phi^{r+1}(\tau)] \\ &\quad = W_{n+1|r}(\bar{\sigma}_{n+1|r}(\tau')\,|\,\tau') - W_{n+1|r}(\bar{\sigma}_{n+1|r}(\tau')\,|\,\tau) \\ &\quad \geq W^*_{n+1|r}(\tau') - W^*_{n+1|r}(\tau) \\ &\quad \geq W_{n+1|r}(\underline{\sigma}_{n+1|r}(\tau)|\tau') - W^*_{n+1|r}(\tau) \\ &\quad \geq \big[1-\underline{\sigma}_{n+1|r}(\tau)\big][\Phi^{r+1}(\tau') - \Phi^{r+1}(\tau)\big]. \end{aligned}$$

This leaves $\underline{\sigma}_{n+1|r}(\tau) \geq \bar{\sigma}_{n+1|r}(\tau')$. It follows that $\gamma_{n+1|r}(\tau)$ is strictly decreasing in τ since $\gamma_{n+1|r}(\tau) = (1-\delta)(q-d) + \delta k_S[1-\tau(1-q)] + \delta(1-k_S)(1-\tau)\hat{\sigma}_{n+1|r}(\tau)$ and $\hat{\sigma}_{n+1|r}(\tau)$ is nonincreasing. Hence, condition (ii) holds for $n+1$.

To establish (iii), consider the state and expected next offer induced by S's offer of y in $W_{n+2|r-1}(y|\hat{\tau})$. By construction, the subsequent cut point τ satisfies $y \in \gamma_{n+1|r}(\tau)$. But $\gamma_{n+1|r}(\tau)$ is strictly decreasing. Hence, a single τ satisfies the condition that $y \in \gamma_{n+1|r}(\tau)$, and this defines τ as a function of y. The expected next offer

e is also a function of y since $y \in \gamma_{n+1|r}(\tau)$ implies $y = (1-\delta)(q-d) + \delta k_S[1-\tau(1-q)] + \delta(1-k_S)(1-\tau)e$ for some $e \in \hat{\sigma}_{n+1|r}(\tau)$ and τ is a function of y. Call these functions $\tilde{\tau}_{n+2|r-1}(y)$ and $e_{n+1|r}(y)$.

To see that $\tilde{\tau}_{n+2|r-1}(y)$ is nonincreasing, assume the contrary. Then there is a $y' > y$ such that $\tau' > \tau$ where $\tau' \equiv \tilde{\tau}_{n+2|r-1}(y')$ and $\tau \equiv \tilde{\tau}_{n+2|r-1}(y)$. Since $y' \in \gamma_{n+1|r}(\tau')$ and $y \in \gamma_{n+1|r}(\tau)$, it follows that $\bar{\gamma}_{n+1|r}(\tau') - \underline{\gamma}_{n+1|r}(\tau) \geq y' - y > 0$ which in turn implies $\delta(1-k_S)[(1-\tau')\bar{\sigma}_{n+1|r}(\tau') - (1-\tau)\underline{\sigma}_{n+1|r}(\tau)] > \delta k_S(1-q)(\tau'-\tau) > 0$. Since $\hat{\sigma}_{n+1|r}(\tau)$ is nonincreasing, the previous inequality reduces to the contradiction $0 > (\tau-\tau')\underline{\sigma}_{n+1|r}(\tau) > 0$. This contradiction guarantees $\tilde{\tau}_{n+2|r-1}(y)$ is nonincreasing. That the graph of $\gamma_{n+1|r}(\tau)$ is connected and strictly decreasing also ensures that $\tilde{\tau}_{n+2|r-1}(y)$ is continuous.

Finally, $e_{n+1|r}(y) = [y - (1-\delta)(q-d) - \delta k_S[1-\tilde{\tau}_{n+2|r-1}(y)(1-q)]]/[\delta(1-k_S)(1-\tilde{\tau}_{n+2|r-1}(y))]$ is clearly continuous since $\tilde{\tau}_{n+2|r-1}(y)$ is. To see that $e_{n+1|r}(y)$ is nondecreasing, suppose $y' > y$. Since $\tilde{\tau}_{n+2|r-1}(y)$ is nonincreasing, $\tau' \leq \tau$ where $\tau' \equiv \tilde{\tau}_{n+2|r-1}(y')$ and $\tau \equiv \tilde{\tau}_{n+2|r-1}(y)$. If $\tau' = \tau$, then the expression for $e_{n+1|r}(y)$ shows that $e_{n+1|r}(y') - e_{n+1|r}(y) = (y'-y)/[\delta(1-k_S)(1-\tau)] > 0$. If $\tau' < \tau$, then $e_{n+1|r}(y') \geq \underline{\sigma}_{n+1|r}(\tau') \geq \bar{\sigma}_{n+1|r}(\tau) \geq e_{n+1|r}(y)$ since $\hat{\sigma}_{n+1|r}(\tau)$ is nonincreasing. Thus, $e_{n+1}(y)$ is nondecreasing and (iii) and holds for $n+1$. ■

Three more lemmas are needed in order to specify the dynamic programing problem. The first shows that being constrained to make no more than n offers does not affect the set of optimal offers as long as n is sufficiently large. (Ultimately this will ensure that S only makes finitely many equilibrium offers.)

Lemma 6. *There exists a finite $\bar{N}$ such that $W^*_{\bar{N}|r}(\tau) = W^*_{\bar{N}+j|r}(\tau)$ and $\sigma_{\bar{N}|r}(\tau) = \sigma_{\bar{N}+j|r}(\tau)$ for all $\tau \in [\underline{\tau}, \bar{\tau}]$, $r \geq 0$, and $j \geq 0$.*

The intuition underlying this result is that each realization or fight shifts the distribution of τ toward $\underline{\tau}$. That is, $\Theta^r(\tau)$ is increasing in r. After a sufficiently large number of realizations, so much of the mass is close to $\underline{\tau}$ that S's optimal offer is $\tilde{x}(\underline{\tau})$. Hence the number of offers is bounded. (Again, proofs are in Powell 2003.)

Now take N to be the smallest $\bar{N}$ such that limiting S to no more than $\bar{N}$ offers before proposing $\tilde{x}(\underline{\tau})$ does not actually constrain S. Lemma 7 demonstrates that the distance between any two states is bounded away from zero:

Lemma 7. *Let (τ_i, m) and $(\tau_{i-1}, m+1)$ be any two consecutive states along any solution to $W^*_{n|r}(\tau)$ where $\tau_i \geq \hat{\tau}$ and $\hat{\tau}$ is as defined in Lemma 4 (i.e., $\sigma_{n|r}(\tau) = \{\tilde{x}(\underline{\tau})\}$ for any $\tau \leq \hat{\tau}$). Then, there exists an $\varepsilon > 0$ such that $\tau_i - \tau_{i-1} > \varepsilon$ for any $n \geq 0$ and $r \geq 0$.*

Finally, observe that if S offers y in $W_{n|r}(y|\tau')$, then Lemma 5 ensures that S's expected next offer e is an element of the convexification of $\sigma_{n-1|r+1}(\tau)$ given by $y \in \gamma_{n-1|r+1}(\tau'')$. If e does not maximize $W_{n-1|r+1}(\tau'')$, i.e., if $e \notin \sigma_{n-1|r+1}(\tau'')$, then the game-theoretic interpretation of e is that it is part of a mixed strategy in which S randomizes over elements of $\sigma_{n-1|r+1}(\tau'')$. This suggests that there may be a substantial amount of mixing along the path that maximizes $W_{n|r}(y|\tau')$. Lemma 8 shows that this is not the case. S may randomize over the elements of $\sigma_{n|r}(\tau')$ when maximizing $W_{n|r}(y|\tau')$. But all of S's subsequent offers are deterministic.

Lemma 8. *If S offers $y \in \sigma_{n|r}(\tau')$ for any $n, r \geq 0$, then S's subsequent offers are deterministic and defined by $y = \underline{\gamma}_{n-1|r+1}(\tau)$.*

The preceding lemmas make it possible to define a dynamic programming problem such that: (i) S can make arbitrarily many offers, and (ii) proposing y in state $(\hat{\tau}, r)$ induces a cut point defined by $y \in \gamma_{N|r}(\tau)$. To this end, let

$$
\begin{aligned}
M_r(\tau', \{y_m\}_{m=0}^{\infty}) \\
\equiv \sum_{m=0}^{\infty} \Big[& \delta^m (1-k_S)^m \big((1-y_m)\big[\Phi^{r+m}\big(\gamma^{-1}_{N|r+m}(y_{m-1})\big) \\
& - \Phi^{r+m}\big(\gamma^{-1}_{N|r+m}(y_m)\big) + [(1-\delta)(1-q-s) \\
& + \delta(1-qk_S)]\Phi^{r+m}\big(\gamma^{-1}_{N|r+m}(y_m)\big) \\
& - \delta(1-qk_S)\Phi^{r+m+1}\big(\gamma^{-1}_{N|r+m}(y_m)\big)\big)\Big]
\end{aligned}
$$

where y_0 is S's initial offer, y_1 the next and so on, and where the dummy y_{-1} is defined as $\underline{\gamma}_{N|r}(\tau')$ so that $\gamma^{-1}_{N|r}(y_{-1}) = \tau'$. Then $M^*_r(\tau') \equiv \max_{\{y_m\} \in \mathcal{S}} M_r(\tau', \{y_m\}_{m=0}^{\infty})$ where $\mathcal{S}$ is the set of sequences $\{y_m\}_{m=0}^{\infty}$ such that $y_{-1} \leq y_m \leq y_{m+1} \leq \tilde{x}(\underline{\tau})$.

Lemma 9. *The solutions to $M^*_r(\tau')$ for $\tau' \in [\underline{\tau}, \bar{\tau}]$ are identical to the solutions of $W^*_{N|r}(\tau')$.*

The Equilibrium

It follows that a sequence of offers and states is a PBE path starting in state (τ', r) if and only if it is a solution to $W^*_{N|r}(\tau')$. Consider the "only if" part first:

Lemma 10. *Any solution to $W^*_{N|r}(\tau')$ is a PBE path starting in state (τ', r).*

To construct an equilibrium assessment based on a solution to $W^*_{N|r}(\tau')$, let the "fighting" history $f_n \equiv \{y_0, \ldots, y_{n-1}\}$ be any sequence of offers in response to which D has always fought, and let $\tau(f_n)$ be the state induced by these offers, i.e., $\tau_n(f_n) \equiv \min\{\bar{\tau}, \gamma^{-1}_{N|r+1}(y_0), \ldots, \gamma^{-1}_{N|r+n}(y_{n-1})\}$. Taking $\{x^*_{n|r+n}, \tau_{n|n+r}\}_{n=0}^{N-1}$ to be any solution to $W^*_{N|r}(\tau')$, define the following strategies and beliefs: Given that the states have fought m times, then type τ accepts x if $x \geq \underline{\gamma}_{N|m+1}(\tau)$ and fights if $x < \underline{\gamma}_{N|m+1}(\tau)$. S's strategy is: (i) offer $x^*_{0|r}$ in the first round; (ii) offer $\underline{\sigma}_{N|r+n}(\tau(f_n))$ following f_n; (iii) offer $\tilde{x}(\bar{\tau})$ following any history along which D waited at least once; and (iv) wait in any around in which D waits. S also believes τ is distributed according to $\Theta^{r+n}(\tau)/\Theta^{r+n}(\tau(f_n))$ after f_n and $\tau = \bar{\tau}$ with probability one if D has ever waited. Verifying that this assessment is a PBE is tedious but straightforward.

Two definitions help demonstrate the "if" part, i.e., that any equilibrium path must be a solution to $W^*_{N|0}(\tau')$. Let (τ_i, r) be the current state following any history in any PBE. Suppose further that S offers y (which need not be an equilibrium offer). Then, D is *playing according to the strategy* $\gamma_{N|0}(\tau)$ if all $\tau > \min\{\gamma^{-1}_{N|r+1}(y), \tau_i\}$ accept y and all $\tau < \min\{\gamma^{-1}_{N|r+1}(y), \tau_i\}$ fight when offered y. Note further that following this strategy implies stationarity. Regardless of the history of previous offers, D always responds in the same way to y. As for the second definition, the pair $(\hat{\tau}, \underline{\gamma}_{N|0}(\tau))$ is a *reservation-offer strategy* if the cut point in the next state, τ_{i+1}, satisfies two conditions for all $\tau' < \min\{\hat{\tau}, \tau_i\}$: (i) $\tau_{i+1} \leq \tau'$ if $y > \underline{\gamma}_{N|r+1}(\tau')$, and, (ii) $\tau_{i+1} > \tau'$ if $y < \underline{\gamma}_{N|r+1}(\tau')$.[21]

Clearly, D is playing according to $\gamma_{N|0}(\tau)$ if $(\bar{\tau}, \underline{\gamma}_{N|0}(\tau))$ is a reservation-offer strategy. And, D's playing according to $\gamma_{N|0}(\tau)$ implies that S's best reply in state (τ', r) must satisfy $M^*_r(\tau')$ or, equivalently, $W^*_{N|r}(\tau')$. Hence, any PBE path must be a solution to $W^*_{N|0}(\tau')$ if $(\bar{\tau}, \underline{\gamma}_{N|0}(\tau))$ is a reservation-offer strategy. Lemma 11 demonstrates this:

Lemma 11. *The pair $(\bar{\tau}, \underline{\gamma}_{N|0}(\tau))$ is a reservation-offer strategy.*

Proof: The argument takes three steps. The first shows that a reservation-offer strategy actually exists for a $\hat{\tau}$. The second demonstrates that there exists an $\varepsilon > 0$ such that if the current state is (τ', r) and if $\tau' \leq \hat{\tau} + \varepsilon$, then S's continuation payoff in any PBE is $M^*_r(\tau')$. The third step establishes that $(\hat{\tau} + \varepsilon, \underline{\gamma}_{N|0}(\tau))$ is an reservation-offer strategy. It then will be evident that the second and third steps can be repeated finitely many times to establish that $(\bar{\tau}, \underline{\gamma}_{N|0}(\tau))$ is reservation-offer strategy.

Lemma 4 guarantees that there exists a $\hat{\tau}_1 > \underline{\tau}$ such that S offers $\tilde{x}(\underline{\tau})$ whenever the state is no more than $\hat{\tau}_1$. It follows that $(\hat{\tau}_1, \underline{\gamma}_{N|0}(\tau))$ is a reservation-offer strategy. To see that (i) holds, suppose S offers y in state (τ_i, r) and $y > \underline{\gamma}_{N|r+1}(\hat{\tau})$ for $\hat{\tau} < \min\{\hat{\tau}_1, \tau'\}$. Arguing by contradiction, assume further that $\tau_{i+1} > \hat{\tau}$. This assumption implies that $\hat{\tau}$ rejects y. However, the best that $\hat{\tau}$ can do by rejecting y is to obtain $\tilde{x}(\underline{\tau})$ after fighting for one round. Hence, the contradiction $y \leq [1 - \delta(1 - k(\hat{\tau}))]\tilde{x}(\hat{\tau}) + \delta(1 - k(\hat{\tau}))\tilde{x}(\underline{\tau}) = \underline{\gamma}_{N|r+1}(\hat{\tau})$ where the equality holds because $\sigma_{N|r+1}(\hat{\tau}) = \{\tilde{x}(\underline{\tau})\}$ since $\hat{\tau} < \hat{\tau}_1$. Thus, (i) holds.

Turning to (ii), assume $y < \underline{\gamma}_{N|r+1}(\hat{\tau})$ and $\tau_{i+1} \leq \hat{\tau}$ where $\hat{\tau} < \min\{\hat{\tau}_1, \tau_i\}$. The latter implies that $\hat{\tau}$ at least weakly prefers y to fighting on. It also implies $\tau_{i+1} < \hat{\tau}_1$, which means that S's offer in state $(\tau_{i+1}, r+1)$ in any PBE is $\tilde{x}(\underline{\tau})$. Consequently, $\hat{\tau}$'s payoff to rejecting y is $[1 - \delta(1 - k(\hat{\tau}))]\tilde{x}(\hat{\tau}) + \delta(1 - k(\hat{\tau}))\tilde{x}(\underline{\tau}) = \underline{\gamma}_{N|r+1}(\hat{\tau})$. But $\hat{\tau}$'s weak preference for y yields the contradiction that $y \geq \underline{\gamma}_{N|r+1}(\hat{\tau})$, and this contradiction ensures (ii) holds. Thus, $(\hat{\tau}_1, \underline{\gamma}_{N|0}(\tau))$ is a reservation-offer strategy.

Now define $R^{\Sigma}_r(\tau')$ to be S's continuation payoff in the PBE Σ (normalized by multiplying $(1 - \delta)\Phi^r(\tau')$) starting in state (τ', r). Then, there exists an $\varepsilon > 0$ such that $R^{\Sigma}{}_r(\tau') = W^*_{N|r}(\tau')$ for all $\tau' \leq \max\{\underline{\tau}, \hat{\tau}_2\}$ where $\hat{\tau}_2 \equiv \hat{\tau}_1 + \varepsilon$. To establish this claim, observe that there is nothing to show if $\tau' \leq \hat{\tau}_1$. In this case, S offers $\tilde{x}(\underline{\tau})$ in any PBE and $R^{\Sigma}{}_r(\tau') = W^*_{N|r}(\tau')$.

The first step in demonstrating the claim for $\tau' > \hat{\tau}_1$ is defining ε. Lemma 7 ensures that there exists an $\varepsilon' > 0$ such that the difference between any two cut points along any solution to $W^*_{N|r}(\tau')$ and therefore $M^*_r(\tau')$ is at least ε'. And, as shown below, there also is a $\xi > 0$ such that $(1 - \delta)V^*_{N|r}(\tau') > \xi + B_r(\tau')$ for all $\tau' \in [\hat{\tau}_1, \bar{\tau}]$ and $r \geq 0$ where $B_r(\tau')$ is S's (average) payoff to making an offer that is sure to be rejected, i.e.,

$$B_r(\tau') \equiv \int_{\underline{\tau}}^{\tau'} \Big[(1 - \delta)(1 - q - s) + \delta k(\tau)(1 - p(\tau)) + \delta(1 - \delta) \times (1 - k(\tau))V^*_{n|r+1}(\tau')\Big] \frac{d\Theta^r(\tau)}{\Theta^r(\tau')}.$$

Define $\varepsilon \equiv \min\{\varepsilon', \xi\Theta(\hat{\tau}_1)/[2(1 - \tilde{x}(\bar{\tau}))\bar{\theta}]\}$.

[21] The definition of a reservation-offer strategy as well as the proof of Lemma 11 draw heavily on Gul, Sonnenschein, and Wilson (1986).

To see that ξ exists note that $(1-\delta)V^*_{N|r}(\tau') > B_r(\tau')$ since it is never optimal to make an offer that is sure to be rejected. The difference $(1-\delta)V^*_{N|r}(\tau') - B_r(\tau')$ is continuous in τ' since $V^*_{N|r}(\tau')$ and $B_r(\tau')$ are continuous, and τ' is an element of the compact set $[\hat{\tau}_1, \bar{\tau}]$. Hence, $(1-\delta)V^*_{N|r}(\tau') - B_r(\tau')$ takes on its minimum value for some τ'' in this set. This means that for each $r \geq 0$, there exists an $\xi_r > 0$ such that $(1-\delta)V^*_{N|r}(\tau') - B_r(\tau') \geq (1-\delta)V^*_{N|r}(\tau'') - B_r(\tau'') > \xi_r$ for all $\tau' \in [\hat{\tau}_1, \bar{\tau}]$.

To show that the ξ_r are bounded away from zero, take r large enough to ensure that $(1-\delta)V^*_{N|r+1}(\tau') < 1 - \tilde{x}(\underline{\tau}) + (1-\delta)(s+d)/4$. To see that this can be done, take $\tau^* < \hat{\tau}_1$ small enough to guarantee that $\tilde{x}(\underline{\tau}) - \tilde{x}(\tau^*) < (1-\delta)(s+d)/4$. Also take R_1 large enough to guarantee that S's initial offer induces a state bounded above by τ^* for all $\tau' \in [\hat{\tau}_1, \bar{\tau}]$ whenever $r \geq R_1$. The proof of Lemma 6 (see Powell 2003) ensures that such an R_1 exists. The basic idea is that every time D fights without collapsing, the mass of the distribution of τ shifts toward $\underline{\tau}$. For r sufficiently large, so much of the mass is so close to $\underline{\tau}$ that S's optimal offer induces a cut point below τ^*.

Then, S's initial offer in $V^*_{N|r+1}(\tau')$, say y, must be at least as large as $\tilde{x}(\tau^*)$ since y induces a cut point bounded above by τ^* and τ^* would never accept anything less than $\tilde{x}(\tau^*)$. Moreover, S's payoff is bounded above by what S receives if this offer is accepted for sure. Hence, $(1-\delta)V^*_{N|r+1}(\tau') \leq 1 - y \leq 1 - \tilde{x}(\tau^*)$. But, τ^* is constructed so that $1 - \tilde{x}(\tau^*) < 1 - \tilde{x}(\underline{\tau}) + (1-\delta)(s+d)/4$.

Substituting this bound for $(1-\delta)V^*_{N|r+1}(\tau')$ in the expression for $B_r(\tau')$ gives

$$B_r(\tau') < 1 - \tilde{x}(\underline{\tau}) - \frac{3(1-\delta)(s+d)}{4} + \int_{\underline{\tau}}^{\tau'} [(1-\delta(1-k(\tau)))(\tilde{x}(\underline{\tau}) - \tilde{x}(\tau))] \times \frac{d\Theta^r(\tau)}{\Theta^r(\tau')}.$$

Since the integrand is bounded above by $\tilde{x}(\underline{\tau}) - \tilde{x}(\tau)$ and $\tilde{x}(\tau)$ is decreasing in τ, the integral is bounded above by

$$[\tilde{x}(\underline{\tau}) - \tilde{x}(\tilde{\tau})]\frac{\Theta^r(\tilde{\tau})}{\Theta^r(\tau')} + [\tilde{x}(\underline{\tau}) - \tilde{x}(\tau')] \times \frac{\Theta^r(\tau') - \Theta^r(\tilde{\tau})}{\Theta^r(\tau')}$$

for any $\tilde{\tau} \in [\underline{\tau}, \tau']$. Now take $\tilde{\tau}$ close enough to $\underline{\tau}$ so that $\tilde{x}(\underline{\tau}) - \tilde{x}(\tilde{\tau}) < (1-\delta)(s+d)/4$, and R_2 large enough to ensure that so much of the mass of $\Theta^r(\tau)$ is so close to $\underline{\tau}$ that the second term is also less than $(1-\delta)(s+d)/4$. This leaves $B_r(\tau') < 1 - \tilde{x}(\underline{\tau}) - (1-\delta)(s+d)/4$ for $\tau' \in [\hat{\tau}_1, \bar{\tau}]$ and $r \geq R \equiv \max\{R_1, R_2\}$. That S can always offer $\tilde{x}(\underline{\tau})$ implies $(1-\delta)V^*_{N|r}(\tau') \geq 1 - \tilde{x}(\underline{\tau})$. Hence there exists a $\xi_R > 0$ such that $(1-\delta)V^*_{N|r}(\tau') > \xi_R + B_r(\tau')$ for all $\tau' \in [\hat{\tau}_1, \bar{\tau}]$ if $r \geq R$. Consequently, the ξ_r are bounded away from zero, and there exists a ξ such that $(1-\delta)V^*_{N|r}(\tau') > \xi + B_r(\tau')$ for all $\tau' \in [\hat{\tau}_1, \bar{\tau}]$ and $r \geq 0$.

Having defined ε, return to the claim that $R^\Sigma_r(\tau') = W^*_{N|r}(\tau')$ for all $\tau' \leq \max\{\underline{\tau}, \hat{\tau}_2\}$ where $\hat{\tau}_2 = \hat{\tau}_1 + \varepsilon$. The first step in verifying this claim is showing $R^\Sigma_r(\tau') \geq W^*_{N|r}(\tau')$. To this end, let $\{y^*_m\}^\infty_{m=0}$ be a solution to $M^*_r(\tau')$ and type $\tau^*_0 = \gamma^{-1}_{N|r}(y^*_0)$ be the state induced by the first offer y^*_0. The construction of ε ensures that the distance between states is greater than ε, so τ^*_0 is strictly less than $\tau' - \varepsilon$. Consequently, $\tau^*_0 < \hat{\tau}_1$, and $y^*_0 \geq \underline{\gamma}_{N|r+1}(\tau^*_0) > \underline{\gamma}_{N|r+1}(\hat{\tau}_1)$. Furthermore, S can obtain at least $M^*_r(\tau')$ in equilibrium by offering y^*_0 because this offer induces state τ^*_0 as $(\hat{\tau}_1, \underline{\gamma}_{N|0}(\tau))$ is a reservation-offer strategy and $\tau^*_0 < \hat{\tau}_1$. Hence, $R^\Sigma_r(\tau') \geq M^*_r(\tau') = W^*_{N|r}(\tau')$.

To establish that $M^*_r(\tau') \geq R^\Sigma_r(\tau')$ and, consequently, that $R^\Sigma_r(\tau') = M^*_r(\tau') = W^*_{N|r}(\tau')$ observe that S would never induce a state $(\tau'', r+1)$ with $\tau'' \geq \hat{\tau}_1$ in equilibrium Σ. Arguing by contradiction to show this, suppose that $\tau'' \geq \hat{\tau}_1$. Then, $R^\Sigma_r(\tau')$ is bounded above by $[1 - \tilde{x}(\tau')][\Phi^r(\tau') - \Phi^r(\hat{\tau}_1)] + [(1-\delta)(1-q-s) + \delta(1-qk_S)]\Phi^r(\tau') - \delta(1-qk_S)\Phi^{r+1}(\tau') + \delta(1-k_S)M^*_{r+1}(\hat{\tau}_1)$. The first term in this expression is an upper bound on S's payoff if $\tau \in [\hat{\tau}_1, \tau']$ weighted by the probability that τ is in this interval. The latter terms are an upper bound on S's continuation payoff given $\tau \leq \tau'$ and that at least one of S's offers has been rejected. (At least one offer must have been rejected in order to have screened out the $\tau > \hat{\tau}_1$.) Observe further that the latter three terms are bounded above by $B_r(\tau')\Phi^r(\tau')$ since $M^*_{r+1}(\hat{\tau}_1) \leq M^*_{r+1}(\tau')$. Hence, $R^\Sigma_r(\tau') \leq [1 - \tilde{x}(\tau')][\Phi^r(\tau') - \Phi^r(\hat{\tau}_1)] + B_r(\tau')\Phi^r(\tau')$. This inequality along with the fact that $M^*_r(\tau') = W^*_{N|r}(\tau') = (1-\delta)V^*_{N|r}(\tau')\Phi^r(\tau') > \xi\ \Phi^r(\tau') + B_r(\tau')\Phi^r(\tau')$ leads to the contradiction $R^\Sigma_r(\tau') < M^*_r(\tau')$ if the definition of ε ensures that $\xi\Phi^r(\tau') > [1 - \tilde{x}(\tau')][\Phi^r(\tau') - \Phi^r(\hat{\tau}_1)]$. And, this contradiction will leave $\tau'' < \hat{\tau}_1$.

To see that ε yields this contradiction, it is enough to show that $\xi > [1 - \tilde{x}(\tau')][1 - \Theta^r(\hat{\tau}_1)/\Theta^r(\tau')]$. Observe first that $\Theta^r(\tau^*)/\Theta^r(\tau')$ is nondecreasing in r for any $\tau^* \leq \tau'$ since the distribution shifts to the left as the number of realizations rises. More precisely,

$$\frac{\Theta^{r+1}(\tau^*)}{\Theta^{r+1}(\tau')} - \frac{\Theta^r(\tau^*)}{\Theta^r(\tau')} = \frac{\Theta^r(\tau^*) - \int_{\underline{\tau}}^{\tau^*} \tau(1-\tau)^r d\Theta}{\Theta^r(\tau') - \int_{\underline{\tau}}^{\tau'} \tau(1-\tau)^r d\Theta} - \frac{\Theta^r(\tau^*)}{\Theta^r(\tau')}$$

$$\geq \frac{\int_{\underline{\tau}}^{\tau^*}(1-\tau)^r d\Theta \int_{\tau^*}^{\tau'} \tau(1-\tau)^r d\Theta - \int_{\tau^*}^{\tau'}(1-\tau)^r d\Theta \int_{\underline{\tau}}^{\tau^*} \tau(1-\tau)^r d\Theta}{\Theta^r(\tau')\left[\Theta^r(\tau') - \int_{\underline{\tau}}^{\tau'} \tau(1-\tau)^r d\Theta\right]}.$$

The expression on the right of the inequality is nonnegative because $\tau^* \int_{\underline{\tau}}^{\tau^*}(1-\tau)^r d\Theta \int_{\tau^*}^{\tau'}(1-\tau)^r d\Theta$ is no larger than the first term in the numerator and no smaller than the second.

That $\Theta^r(\tau^*)/\Theta^r(\tau')$ is nondecreasing implies $[1-\tilde{x}(\tau')][1-\Theta^r(\hat{\tau}_1)/\Theta^r(\tau')] \leq [1-\tilde{x}(\bar{\tau})][1-\Theta(\hat{\tau}_1)/\Theta(\tau')] \leq [1-\tilde{x}(\bar{\tau})]\bar{\theta}(\tau'-\hat{\tau}_1)/\Theta(\hat{\tau}_1)]$. So, $[1-\tilde{x}(\tau')] \times [\Phi^r(\tau') - \Phi^r(\hat{\tau}_1)] < \xi\Phi^r(\tau')$ since, by construction, $\tau' - \hat{\tau}_1 \leq \varepsilon$ and $\varepsilon \leq \xi\Theta(\hat{\tau}_1)/[2\bar{\theta}(1-\tilde{x}(\bar{\tau}))]$. Hence, S would never induce a state $(\tau'', r+1)$ with $\tau'' \geq \hat{\tau}_1$ in Σ.

Combining the facts that S's offer in state τ', say y', induces a subsequent state, τ'', strictly less than $\hat{\tau}_1$ and that $(\hat{\tau}_1, \underline{\gamma}_{N|0}(\tau))$ is a reservation-offer strategy implies that S's payoff to proposing y' is

$$\begin{aligned} R_r^{\Sigma}(\tau') &= (1-y')[\Phi^r(\tau') - \Phi^r(\tau'')] \\ &\quad + [(1-\delta)(1-q-s) + \delta(1-qk_S)]\Phi^r(\tau'') \\ &\quad - \delta(1-qk_S)\Phi^{r+1}(\tau'') + \delta(1-k_S)M_{r+1}^*(\tau'') \end{aligned} \tag{A2}$$

These facts also imply that if y'' is S's possibly mixed equilibrium offer in state $(\tau'', r+1)$, then the support of y'' are elements of $\sigma_{N|r+1}(\tau'')$ since all PBE paths of games in states $(\tau'', r+1)$ with $\tau'' < \hat{\tau}_1$ must be solutions to $W_{N|r+1}^*(\tau'')$. Thus, $y'' \geq \underline{\sigma}_{N|r+1}(\tau'')$. It must also be that $y' \geq \underline{\gamma}_{N|r+1}(\tau'')$. Otherwise a neighborhood of types slightly greater than τ'' would have preferred fighting for y'' to accepting y'. But this would contradict the fact that y induces τ''.

Now consider S's payoff to offering $\underline{\gamma}_{N|r+1}(\tau'')$ in the program $M_r(\tau', \{y_i\}_{i=0}^{\infty})$ and then playing optimally thereafter. Since $\underline{\gamma}_{N|r+1}(\tau'')$ may not be an optimal offer in state (τ', r), $M_r^*(\tau')$ is at least as large S's payoff to proposing $\underline{\gamma}_{N|r+1}(\tau'')$. Using the Bellman equation to express the latter payoff gives:

$$\begin{aligned} M_r^*(\tau') &\geq [1-\underline{\gamma}_{N|r}(\tau'')]\left[\Phi^r(\tau') - \Phi^r\left(\gamma_{N|r}^{-1}(\underline{\gamma}_{N|r}(\tau''))\right)\right] \\ &\quad + [(1-\delta)(1-q-s) \\ &\quad + \delta(1-qk_S)]\Phi^r\left(\gamma_{N|r}^{-1}(\underline{\gamma}_{N|r}(\tau''))\right) \\ &\quad - \delta(1-qk_S)\Phi^{r+1}\left(\gamma_{N|r}^{-1}(\underline{\gamma}_{N|r}(\tau''))\right) \\ &\quad + \delta(1-k_S)M_{r+1}^*\left(\gamma_{N|r}^{-1}(\underline{\gamma}_{N|r}(\tau''))\right) \end{aligned}$$

Since $\gamma_{N|r}^{-1}(\underline{\gamma}_{N|r}(\tau'')) = \tau''$ and $y' \geq \underline{\gamma}_{N|r+1}(\tau'')$, the expression on the right of the previous inequality is at least as large as $R_r^{\Sigma}(\tau')$ (see equation A2 above). Hence, $W_{N|r}^*(\tau') = M_r^*(\tau') \geq R_r^{\Sigma}(\tau')$ which leaves $R_r^{\Sigma}(\tau') = M_r^*(\tau') = W_{N|r}^*(\tau')$.

It follows that $(\hat{\tau}_2, \underline{\gamma}_{N|0}(\tau))$ is a reservation-offer strategy. To see that (i) holds, suppose S offers y in state (τ_i, r), and consider any $\hat{\tau} \in [\hat{\tau}_1, \hat{\tau}_2)$ for which $y > \underline{\gamma}_{N|r+1}(\hat{\tau})$. (If $\hat{\tau} < \hat{\tau}_1$, there is nothing to show as $(\hat{\tau}_1, \underline{\gamma}_{N|0}(\tau))$ is a reservation-offer strategy and consequently $y > \underline{\gamma}_{N|r+1}(\hat{\tau})$ implies $\tau_{i+1} \leq \hat{\tau}$.) Suppose further that the next state, τ_{i+1}, is weakly bounded above by $\hat{\tau}_2$. Then the definition of $\hat{\tau}_2$ ensures that the next state τ_{i+2} is strictly bounded above by $\hat{\tau}_1$. Arguing now by contradiction, assume $\tau_{i+1} > \hat{\tau}$. Then, $\tau_{i+2} < \hat{\tau}_1 \leq \hat{\tau} < \tau_{i+1} \leq \hat{\tau}_2$. Hence, $\hat{\tau}$ accepts in period $i+1$.

Now let y' be in the support of S's possibly mixed offer e' made in state $(\tau_{i+1}, r+1)$. Then $y' \in \sigma_{N|r+1}(\tau_{i+1})$ as $R_{r+1}^{\Sigma}(\tau_{i+1}) = M_{r+1}^*(\tau_{i+1})$ since $\tau_{i+1} \leq \hat{\tau}_2$. This implies e' is an element of $\hat{\sigma}_{N|r+1}(\tau_{i+1})$ and bounded above by $\bar{\sigma}_{N|r+1}(\tau_{i+1})$.

A contradiction follows as $y > \underline{\gamma}_{N|r+1}(\hat{\tau}) = [1 - \delta(1-k(\hat{\tau}))]\tilde{x}(\hat{\tau}) + \delta(1-k(\hat{\tau}))\underline{\sigma}_{N|r+1}(\hat{\tau})$. But $\underline{\sigma}_{N|r+1}(\hat{\tau}) \geq \bar{\sigma}_{N|r+1}(\tau_{i+1})$ since $\tau_{i+1} > \hat{\tau}$ and $\hat{\sigma}_{N|r+1}$ is nonincreasing. This gives $y > [1-\delta(1-k(\hat{\tau}))]\tilde{x}(\hat{\tau}) + \delta(1-k(\hat{\tau}))\bar{\sigma}_{N|r+1}(\tau_{i+1}) \geq [1-\delta(1-k(\hat{\tau}))]\tilde{x}(\hat{\tau}) + \delta(1-k(\hat{\tau}))e'$. Type $\hat{\tau}$, therefore, strictly prefers accepting y in period i to fighting for e' in period $i+1$, and this contradicts the fact that $\hat{\tau}$ accepts in period $i+1$. Thus, (i) holds if $\tau_{i+1} \leq \hat{\tau}_2$.

Now suppose $\tau_{i+1} > \hat{\tau}_2$, and consider that case in which S offers y_{j+1} in state τ_j where $\tau_j > \hat{\tau}_2$ and $\tau_{j+1} \leq \hat{\tau}_2$. (Such an offer must exist as S only makes finitely many proposals before offering $\tilde{x}(\underline{\tau})$.) Since $\tau_j > \tau_{j+1}$, y_{j+1} is accepted with positive probability. Hence, $y_{j+1} > y$ where, recall, y is what S offered in state τ_i. Repeating the argument above again yields a contradiction. Hence, (i) holds for $(\hat{\tau}_2, \underline{\gamma}_{N|r+1}(\tau))$.

Turning to (ii), let τ_i and τ_{i+1} be the current and subsequent cut points. Take a $\hat{\tau}$ and y such that $\hat{\tau} < \min\{\hat{\tau}_2, \tau_i\}$ and $y < \underline{\gamma}_{N|r+1}(\hat{\tau})$. To see that $\tau_{i+1} > \hat{\tau}$, assume the contrary. That $\tau_{i+1} \leq \hat{\tau} < \hat{\tau}_2$ implies $R_{r+1}^{\Sigma}(\tau_{i+1}) = M_{r+1}^*(\tau_{i+1})$ and, therefore, that S offers a

$y' \in \sigma_{N|r+1}(\tau_{i+1})$ in state τ_{i+1}. Hence, S's expected offer in state τ_{i+1}, e', satisfies $\underline{\sigma}_{N|r+1}(\hat{\tau}) \le \underline{\sigma}_{N|r+1}(\tau_{i+1}) \le e'$. This leaves $y < \underline{\gamma}_{N|r+1}(\hat{\tau}) < [1-\delta(1-k(\hat{\tau}))]\tilde{x}(\hat{\tau}) + \delta(1-k(\hat{\tau}))e'$. Since this inequality is strict, there exists a $\zeta > 0$ such that $\hat{\tau} + \zeta < \tau_i$ and all $\tau \in (\hat{\tau}, \hat{\tau} + \zeta)$ strictly prefer fighting for e' to accepting y. But, this contradicts the assumption that τ_{i+1} is the state in period $i+1$ because $\tau_{i+1} \le \hat{\tau} < \hat{\tau} + \zeta < \tau_i$ implies that $\tau \in (\hat{\tau}, \hat{\tau} + \xi)$ are still in the game in period i and prefer not to settle in period i. This contradiction means (ii) holds.

Thus, $(\hat{\tau}_2, \underline{\gamma}_{N|0}(\tau))$ is a reservation-offer strategy. Repeating the argument shows that $(\hat{\tau}_n, \underline{\gamma}_{N|0}(\tau))$ is a reservation-offer strategy where $\hat{\tau}_n = \min\{\hat{\tau}_{n-1} + \varepsilon, \bar{\tau}\}$. Since $\varepsilon > 0$, further repetition eventually ensures that $(\bar{\tau}, \underline{\gamma}_{N|0}(\tau))$ is a reservation-offer strategy. ■

In sum, Lemmas 10 and 11 show that a set of offers and states is a PBE path of the game if and only if it solves $W_{N|0}(\bar{\tau})$. Lemmas 5, 6, and 8 ensure that S's offers are finite, strictly decreasing, and deterministic. This yields the stronger version of Proposition 2:

Proposition 2a. *Suppose that S is satisfied, D is dissatisfied, and that S's beliefs about D's probability of collapse satisfy the technical assumptions discussed above. Then, there exists a smallest $N \ge 0$ such that $\sigma_{N|0}(\tau) = \sigma_{N+j|0}(\tau)$ for all $j \ge 0$ and $\tau \in [\underline{\tau}, \bar{\tau}]$. N defines the set of PBEs: S's initial offer in any PBE is an element of $\sigma_{N|0}(\bar{\tau})$. S's subsequent offers are a function of S's initial offer, strictly decreasing, and deterministic*.

Multiple Offers Per Battle

Turning to the m-offer-per-battle game, define $F_m^D(i)$ for $1 \le i \le m$ to be D's expected payoff to fighting to the finish given i offers to go before the first battle. Then D's payoff to fighting to the finish from the outset of the game when there are m offers before the first battle is $F_m^D(m) = \sum_{n=1}^{\infty}(1-k(\tau))^{n-1}k(\tau)[\sum_{j=0}^{nm-1}(q_m - d_m)\delta_m^j + \delta_m^{nm}[k_S(1-\tau)v_m + k_S\tau q_m]/[(1-\delta_m)[1-(1-\tau)(1-k_S)]]$. If there are fewer than m offers until the first battle ($i < m$), then D rejects i offers before fighting and then rejects m offers before fighting again. Consequently, D's expected payoff is $F_m^D(i) = \sum_{j=0}^{i-1}(q_m - d_m)\delta_m^j + \delta_m^i k(\kappa)p(\tau)v_m/(1-\delta_m) + (1-k(\tau))\delta_m^i F_m^D(m)$.

A type is dissatisfied in the m-offer game as in the 1-offer game, if its expected payoff to fighting to the finish from the outset is strictly larger than the payoff to living with the status quo. Thus, d_m is dissatisfied if $F_m^D(m) > q_m/(1-\delta_m)$. Rewriting this expression in terms of q, d, and δ shows that D is dissatisfied in the m-offer game if and only if it is dissatisfied in the 1-offer game, i.e., if and only if $\delta k(p-q) - d(1-\delta) > 0$. Note that the closer a dissatisfied type is to fighting, the higher its payoff to fighting to the finish: $F_m^D(i') > F_m^D(i)$ for $1 \le i' < i \le m$.

Now consider an m-offer game in which S is uncertain of D's cost but believes it to be distributed according to $G_m(d_m) \equiv G(d_m(1-\delta)/(1-\delta_m))$. If m is sufficiently large, agreement is reached almost immediately with no fighting and Proposition 3 as stated above holds. The proof is a very tedious extension of a standard argument. (See Fudenberg and Tirole (1999, 411-12) for a discussion.) The basic idea is to show that if Proposition 3 does not hold, then S can increase its payoff by skipping an offer and speeding up the bargaining.

Proposition 4 shows that if S is uncertain of the distribution of power, then there is a positive probability of fighting whenever there is any bargaining. The proposition also establishes a condition sufficient to ensure that there is some delay.

Proposition 4. *Suppose that S is uncertain about the distribution of power in the m-offer-per-battle game. In equilibrium, either S initially offers $\tilde{x}(\underline{\tau})$ and the game ends immediately, or $\underline{\tau}$ strictly prefers to wait until after the first battle to accept an offer in which case there is a positive probability that the states fight at least one battle. Suppose, further, that $x^* \in \sigma_{1|0}(\bar{\tau})$; that $\tilde{\tau}(x^*)$ is the type that is indifferent between accepting x^* and fighting one battle for $\tilde{x}(\underline{\tau})$ in the one-offer-per-battle game; and that $[(\tilde{x}(\underline{\tau}) - x^*)/(1-\delta)](1-\Theta(\tilde{\tau}(x^*)) > s + d$ for any $x^* \in \sigma_{1|0}(\bar{\tau})$. Then there is always delay in equilibrium.*

Proof: The argument for the first claim is sketched above in the text. The key to the second claim is to consider an x_m that if offered on the eve of battle in the m-offer game induces the same cut point that $x^* \in \arg\max V_{1|0}(\bar{\tau})$ would induce in a one-offer-per-battle game. ■

References

Admati, Anat, and Motty Perry. 1987. "Strategic Delay in Bargaining." *Review of Economic Studies* 54(July):345–64.

Ausubel, Lawrence, and Raymond Deneckere. 1989. "Reputation in Bargaining and Durable Goods Monopoly." *Econometrica* 57 (May):511–31.

Blainey, Geoffrey. 1973. *The Causes of War*. New York: Free Press.

Fearon, James. 1994a. "Domestic Political Audiences and the Escalation of International Disputes." *American Political Science Review* (September):577–92.

Fearon, James. 1994b. "Signaling Versus the Balance of Power and Interests." *Journal of Conflict Resolution* (February):236–69.

Fearon, James. 1995. "Rationalist Explanations for War." *International Organization* (Summer):379–414.

Filson, Darren, and Suzanne Werner. 2002a. "A Bargaining Model of War and Peace." *American Journal of Political Science* 46(October):819–38.

Filson, Darren, and Suzanne Werner. 2002b. "The Dynamics of Bargaining in War." Manuscript, Department of Political Science, Emory University.

Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory*. Cambridge, MA: MIT Press.

Goemans, Hein. 2000. *War and Punishment*. Princeton: Princeton University Press.

Gul, Faruk, and Hugo Sonnenschein. 1988. "On Delay in Bargaining with One-Sided Uncertainty." *Econometrica* 56(May):601–11.

Gul, Faruk, Hugo Sonnenschein, and Robert Wilson. 1986. "Foundations of Dynamic Monopoly and the Coase Conjecture." *Journal of Economic Theory* 39(June):155–90.

Heifetz, Avid, and Ella Segev. (2002). "Escalation and Delay in Long International Conflicts." Manuscript, School of Economics, Tel Aviv University.

Hickey, Donald. 1989. *The War of 1812*. Chicago: University of Illinois Press.

Kim, Dong-Won. 2002. "Bargaining and War Initiation." Manuscript, Department of Political Science, Texas Tech University.

Kydd, Andrew. 1997. "Game Theory and the Spiral Model." *World Politics* 49(April):371–400.

Kydd, Andrew. 2001. "Which Side Are You On? Mediation as Cheap Talk." *American Journal of Political Science* 47(October): 597–611.

Morris, Stephen. 1995. "The Common Prior Assumption in Economic Theory." *Economics and Philosophy* 11(October): 227–53.

Morrow, James D. 1989. "Capabilities, Uncertainty, and Resolve." *American Journal of Political Science* 33(November): 941–72.

Muthoo, Abhinay. 1999. *Bargaining Theory with Applications*. New York: Cambridge University Press.

Powell, Robert. 1996a. "Bargaining in the Shadow of Power." *Games and Economic Behavior* 15(August):255–89.

Powell, Robert. 1996b. "Stability and the Distribution of Power." *World Politics* 48(January):239–67.

Powell, Robert. 1999. *In the Shadow of Power*. Princeton: Princeton University Press.

Powell, Robert. 2002. "Bargaining Theory and International Conflict." *Annual Review of Political Science* 5(June):1–30.

Powell, Robert. 2003. "Bargaining and Learning While Fighting: Proofs." Manuscript, Department of Political Science, UC Berkeley.

Schultz, Kenneth. 1999. "Do Democratic Institutions Constrain or Inform." *International Organization* 53(Spring):23–66.

Slantchev, Branislav. 2002. "The Power to Hurt: Costly Conflict with Completely Informed States." *American Political Science Review* 47(February):1–11.

Slantchev, Branislav. 2003. "The Principle of Convergence in Wartime Negotiations." *American Political Science Review*. Forthcoming.

Smith, Alastair. 1998a. "Fighting Battles, Winning Wars." *Journal of Conflict Resolution* 42(June):301–20.

Smith, Alastair. 1998b. "International Crises and Domestic Politics." *American Political Science Review* 92(September):623–38.

Smith, Alastair, and Alan Stam. 2003. "Bargaining through Conflict." Manuscript, Department of Politics, New York University.

Wagner, R. Harrison. 1994. "Peace, War, and the Balance of Power." *American Political Science Review* 88(September): 593–607.

Wagner, R. Harrison. 2000. "Bargaining and War." *American Journal of Political Science* 44(July):469–84.

Werner, Suzanne. 2000. "Deterring Intervention." *American Journal of Political Science* 46(October):720–32.

Wittman, Donald. 1979. "How a War Ends." *Journal of Conflict Resolution* 23(December):743–63.