UNSW Business School Research Paper No. 2015 ECON 07
First version April 2015
UNSW Business School Research Paper No. 2015 ECON 07A
Second version February 2017

Intimidation: Linking Negotiation and Conflict

Sambuddha Ghosh
Gabriele Gratton
Caixia Shen

Intimidation:
Linking Negotiation and Conflict

Sambuddha Ghosh, Gabriele Gratton, and Caixia Shen*

February 14, 2017

Abstract

A challenger wants a resource initially held by a defender, who can negotiate a settlement by offering to share the resource; if this fails, conflict ensues. We propose a dynamic non-cooperative game in which negotiation and conflict influence each other. During conflict each player could be a tough type for whom fighting is costless; therefore non-concession intimidates the opponent into conceding. We derive how the probabilities of being tough, relative discount factors, and the time until conflict all determine the likelihood of conflict and its length. Several policy relevant lessons emerge. An offer made by a defender unaware of his own type fails with some probability if the threat of conflict is imminent, but succeeds for sure if the threat of conflict is distant and the defender is more patient than the challenger. Instead, if the offer is made by a privately informed and more patient defender, negotiation fails completely. Finally, allowing multiple offers is no more effective---a serious offer is made only on the brink of conflict. Conflict can be prolonged even if the probability of the tough type is arbitrarily small.

Keywords: Intimidation, reputation, terrorism, negotiation, brinkmanship, costly war-of-attrition. **JEL Classification Numbers:** D74, D82.

---

*Ghosh: School of Economics, Shanghai University of Finance and Economics ,777 Guoding Road, Shanghai, China 200433, sghosh@mail.shufe.edu.cn. Gratton: School of Economics, UNSW Business School, UNSW Sydney, NSW 2052, Australia, g.gratton@unsw.edu.au. Shen: School of International Business Administration, Shanghai University of Finance and Economics, 777 Guoding Road, Shanghai, China 200433, shencaixia@gmail.com. We are grateful to Masaki Aoyagi, Gaurav Aryal, Laurent Bouton, Alessandra Casella, Micael Castanheira, Costas Cavounidis, Pauline Grosjean, Anton Kolotilin, Hongyi Li, Qingmin Liu, Youming Liu, Massimo Morelli, Santiago Oliveros, Juan Ortner, Carlos Pimienta, Ronny Razin, Santiago Sanchez-Pages, Balázs Szentes, John Tang, and Adam Wong. We thank seminar participants at the Australian National University, Boston University, Columbia University, European University Institute, Georgetown University, University of Montreal, the 2014 ASSA Meetings, the 2014 Econometric Society Australasian Meetings, the 2013 Australasian Economic Theory Workshop, and the 2014 SAET Conference.

1 Introduction

Violent conflicts are long, costly, and often impossible to avoid. The latter fact is particularly puzzling: why do negotiations fail to prevent conflict even when compromise solutions are available, commitment is possible, and conflict is likely to be long and painful?[1] Recent work sees uncertainty about the eventual outcome (or costs) of conflict as its fundamental cause. These analyses show that the expected outcome of conflict determines bargaining positions in negotiations (e.g. Esteban and Sàkovics, 2008; Powell, 2004). However, informal discussion also suggests that actions during negotiations affect the expected outcome of conflict, as one party's tough bargaining position intimidates her opponent. We present a model where negotiation and conflict are interlinked. Our unified framework captures several empirical regularities of armed conflics and sheds light on when and why negotiations more likely fail, whether multi-round negotiations lower the chance of eventual conflict, and why conflict can last even though no party gains from it.

Our model has two players, Challenger (she) and Defender (he), each of whom wants a resource that yields flow utility. Defender, who initially holds the resource, can try to negotiate peace by offering a share to Challenger. If the offer is rejected, then Challenger has the option to repeatedly attack Defender and convince him to concede the resource. Our model of conflict is built on the concept of *intimidation*: repeatedly inflicting losses on one's opponent to scare him that further losses will be incurred unless a concession is made. For example, terrorists carry out attacks to raise the specter of further attacks until and unless their demands are met; similarly, workers go on strike to raise the fear of further strikes.[2] However attacking could be coslty too. What allows room for intimidation to function is that conflict is a game of two-sided incomplete information—with an arbitrarily small probability either party to the dispute is a tough type, who does not experience the disutility of conflict. This small uncertainty is magnified by equilibrium play into a significant force that protracts conflict.

The same idea of intimidation also explains why negotiations fail: even if Defender can make offers to Challenger before conflict begins, equilibrium offers are rejected with positive probability by normal (i.e., not tough) Challenger. One might imagine that negotiations fail because Defender is afraid to reveal his type. We show that this intuition is

[1]In a seminal paper, Fearon (1995) poses this as a "central puzzle" that rationalist explanations fail to solve (see also Reiter, 2003). The key word 'rationalist' rules out explanations where the parties to the conflict are entirely or largely irrational.

[2]The analogy between strikes and wars dates back at least to Waltz (1979, p. 114).

incomplete by first looking at the case of Defender being *uninfoirmed*, in the sense of not knowing his cost of pursuing conflict when he makes an offer. Even in this case where offers do not reveal his type, the opportunity to make an offer is a double-edged sword for Defender. On one hand, offers that have a higher probability of being accepted increase the utility of Defender. Indeed, if beliefs were held fixed both before and after the offer, Defender could completely avoid conflict with normal Challenger by offering slightly more than her expected value of entering conflict. On the other hand, if such an offer were made and accepted with certainty in equilibrium, Challenger would prefer to reject it. Indeed, such a rejection would signal that Challenger is tough, thus intimidating Defender into conceding as soon as conflict began. We characterize the sufficient and necessary conditions for the detrimental effect to dominate. Negotiations are more likely to fail if Defender (Challenger) is more (less) patient, is less (more) likely to be tough, and pays a lower (higher) cost pursuing conflict. Furthermore, if the time interval between negotiation and conflict is sufficiently small, conflict begins with strictly positive probability independently of which player is more patient or more likely to be tough.

Let us now turn to multi-round negotiations. We show that in equilibrium all offers, except the one made in the last round, must be unacceptable (no Challenger would ever accept them). Indeed, suppose that before the last round Defender makes an offer that normal Challenger accepts with positive probability. Then Challenger strictly prefers to reject it, as this signals she is tough, thus intimidating Defender into making an even more generous offer in the last round. In other words, offers that could be successful (but weren't) make subsequent offers more costly. Therefore, long negotiations resolve in *brinkmanship*: the parties make no progress towards a peaceful solution up until the last opportunity to negotiate. In particular, as discussed later, the Coase conjecture fails in our setting.

Two broad policy messages emerge. First, both Challenger and Defender benefit from the opportunity to negotiate in two ways, as does any third party interested in mitigating conflict: if negotiation succeeds, then conflict is avoided; if it fails, conflict is deemed to be shorter. Second, when a deadline for making offers exists, multiple rounds of negotiations do not help to avert conflict any more than a single round.

We also study the case of the *informed* Defender, who knows his type at the time he makes offers. We obtain clear-cut conditions under which we have conflict. The key determinant is the ratios of the discount factors between the time of negotiation and the time conflict begins. In contrast with the uninformed Defender case, when Defender is informed about his type, negotiations fail with certainty whenever Challenger is slightly more patient than Defender.

We find that long conflicts can occur even when irrationality becomes vanishingly small, i.e. both players are normal with probability arbitrarily close to 1. Our model shows that conflict extends for a maximum number of periods n, which depends on the probability that the players are tough. When irrationality becomes vanishingly small, under suitable conditions, n tends to infinity. We also explicitly derive the probability of survival of conflict for any period until n. Crucially, the probability that conflict extends to the next period does not depend either on the probabilities that the players are tough or on how long the conflict has been fought on. Thus, our model predicts that conflict is non-duration dependent: the past duration of conflict does not predict its probability to end. In contrast, Reiter (2003) notes that existing models of conflict based on informational asymmetries fail to capture that conflict is in reality non-duration dependent, as these asymmetries should fade in time.

We use our model to make a variety of comparative predictions regarding the likelihood and length of the conflict if negotiations fail. We divide them into two groups. Suppose that a conflict is in progress, i.e. Challenger has attacked at least once and Defender has not conceded yet. The first set of predictions pertain to what happens next. A longer conflict is to be expected if either Challenger or Defender value the resource more, the costs of fighting are smaller, or players are more patient. Most importantly, only these parameters affect the probability that the conflict continues.

Our second set of predictions answer two questions: (i) what is the probability of a first attack? (ii) if Challenger attacks, what is the probability that the Defender concedes before any further attack? To answer these questions, we introduce a notion of what it means for one player to be more committed than the other. In equilibrium, when Challenger is more committed, she attacks for sure. Conversely, when Defender is at least as committed as Challenger, then Challenger mixes between conceding and attacking. Later we show how these predictions help explain some stylized facts about the determinants of interstate war duration, terrorist conflicts, and workers' strikes.

Related Literature

Our main contribution is a simple model that captures the interaction between negotiation and conflict. Fearon (1995) and Powell (1996) study non-cooperative bargaining in the shadow of conflict, treating war as a game-ending move.[3] Powell (2004) and Fearon

[3]Esteban and Sàkovics (2008) consider a similar scenario in a cooperative setting.

(2013) focus on alternating phases of bargaining and conflict where each conflict phase leads to the collapse of one party with an exogenously given probability. In contrast, we model conflict as a *costly* war of attrition driven by reputation *à la* Kreps and Wilson (1982) and Milgrom and Roberts (1982).[4] Thus, in our model violent conflict terminates endogenously with probabilities that are determined in equilibrium, rather than exogenously as in the above models. Among other things, this captures the fact that conflicts last longer when more is at stake. Ponsatì and Sàkovics (1995) and Abreu and Gul (2000) study continuous-time wars of attrition under incomplete information.[5] Our model of conflict uses these insights, but modifies the model to allow for discrete delays, sequential moves, and flow costs and benefits, as these are all natural features in many applications.[6] We also derive conditions under which prolonged conflict arises even when the prior probability of irrationality vanishes.[7] Additionally, we introduce the possibility of negotiation, which is not costly like conflict and can effect compromise outcomes rather than only extreme ones; this allows us to study how negotiation and conflict influence each other.

Lapan and Sandler (1988) model terrorism as a repeated game between players who are irrational with some probability. In their model, absent a concession, the probability of being a commitment type jumps up to an arbitrary and exogenously given quantity. Hodler and Rohner (2012) make this endogenous, but they have only two periods; this in turn means that they predict attacks only when the probability of the terrorist being tough is very large. Our model endogenously determines both the termination of the war of attrition and the evolution of beliefs about the degree of irrationality of one's opponent, and shows that prolonged conflict is compatible with very small degrees of irrationality.

[4]The genesis of the extensive literature on the war of attrition is Maynard-Smith (1974). Hendricks *et al.* (1988) study a *complete-information* war of attrition. Kornhauser *et al.* (1989) apply reputational techniques to a war of attrition in discrete time, with committed types conceding with a constant probability in each period; the main goal is a set of conditions under which the less patient player concedes immediately. In Ordover and Rubinstein (1986) two players fight a war of attrition with a given *finite* deadline; only one player privately knows who will 'win' if the game doesn't end by the deadline; either the uninformed player concedes at the start or players remain indifferent between conceding and persisting. Hörner and Sahuguet (2010) model a war of attrition with two-sided asymmetric information about the *valuation* of the good to each player; when both players can signal their types by sinking costly bids, conflict never goes beyond two periods.

[5]The latter studies a bargaining model where players are, or pretend to be, irrationally committed to mutually incompatible demands. It also shows that as the time interval between offers goes to zero all equilibria of the discrete-time game converge to the unique equilibrium in continuous time.

[6]The conflict part of our model is also similar to the entry deterence model that Kreps and Wilson (1982) touches on, without deriving the equilibrium rigorously. In Chatterjee and Samuelson (1988) both a buyer and a seller have private information and alternately make offers. Although flow utilities are absent, the qualitative features are shared: the soft type mimics the hard type initially. Uniqueness does not hold in their richer setting.

[7]See Fearon (2013) for a brief review of why most bargaining models fail to match the observation that conflicts are usually long.

We share our interest in the length of conflicts with a number of papers. In Brito and Intriligator (1985), and the closely related Sobel and Takahashi (1983), the *uninformed* party uses active conflict to screen among various types of informed parties, who wish to appear stronger than they really are in order to secure a better bargain. Sanchez-Pages (2009a) shows the converse—the *informed* party uses 'limited conflict' to convey credible information to an uninformed party about the eventual outcome of rejecting agreements and triggering 'absolute conflict'; this makes the latter more amenable to agreements even if he were optimistic enough initially to render agreements infeasible. Similarly, in Heifetz and Segev (2005) a party delays making an acceptable offer to credibly signal its true stand, and escalation makes resolution more attractive. See Baliga and Sjöström (2013) and Sanchez-Pages (2009b) for a highly integrated view of the above literature and more.

The literature on bargaining under incomplete information can be traced back to Chatterjee and Samuelson (1983). Fudenberg, Levine, and Tirole (1985) and Gul, Sonnenschein, and Wilson (1986) present dynamic models where a monopolist makes repeated offers to a buyer with unknown valuation. The main result is a verification of the Coase conjecture: as offers become frequent, the monopolist immediately offers a price that is acceptable to the lowest valuation buyer.[8] This is not true in our setting. In fact, our brinkmanship result says that in our setup increasing the number of offers does not ensure that negotiations succeed. Thus, inefficiency survives multiple offers, and is sometimes independent of the exact number.

An immediate implication of our results is that Defender wants to be perceived as tough. Schelling (1956, 1960, 1966) and Crawford (1982) developed the idea that bargaining parties can benefit if they convince their opponent that they are committed to their threat—hence the argument that governments should appear committed to hawkish positions when facing a terrorist threat. Yet, the advantage of showing commitment should not be overstated: once conflict begins, the expected payoff for (normal) Defender is independent of his probability of being tough. In fact, the entire advantage of being perceived as tough comes from the ability to induce a normal Challenger to attack with very low probability. But if the Challenger attacks nonetheless, then Defender must update his beliefs to assign a very high probability to Challenger being tough.

Our idea of intimidation is also related to Silverman (2004), a random-matching model where violence is instrumental in deterring future violence against oneself. If the fraction of agents who directly gain from violence is sufficiently large, then other agents can also engage in it to acquire a reputation for toughness. Yared (2009) considers a defender

[8] Cho (1990) extends the above results to two-sided incomplete information.

with private knowledge of his cost of conceding the flow resource in each period; in equilibrium the challenger attacks with positive probability when no concession is made, so that the defender has an incentive to concede often enough. Since costs are drawn independently across periods, there is no reputation at play, unlike our model.

2 Benchmark Model

In our benchmark model, the parties have a single chance to reach an agreement before conflict begins. In Section 4 we discuss how our results generalize to multiple rounds of negotiation.

Time is continuous and runs from $-T$ to infinity, where T is a fixed positive number. There are two players: Challenger and Defender. They contest a resource of size 1, which is initially held by Defender. Holding the resource gives a flow rent $\tilde{d}$ to Defender and $\tilde{c}$ to Challenger.

The game played by Challenger and Defender is best described by dividing it into two phases: negotiation and conflict. Negotiation runs from $t = -T$ to $t = 0$; conflict runs from $t = 0$ onward.

Negotiation At $t = -T$, Defender can offer a fraction x of the resources to Challenger. Upon observing the offer, Challenger decides to accept or reject the current offer. If Challenger accepts the offer, then the game ends and Challenger and Defender enjoy flow rents $x\tilde{c}$ and $(1 - x)\,\tilde{d}$ thereafter. Otherwise, Defender enjoys the entire resource until time $t = 0$, when the game enters the conflict phase. Thus we interpret T as the time for which Challenger has to wait in order to initiate conflict.[9]

Conflict In the conflict phase, the following two-stage game is played out at each $t \in \{0, \Delta, 2\Delta, \dots\}$, where $\Delta > 0$ is the length of each conflict period:

Stage 1: Challenger chooses whether to fight (attack) or concede;

Stage 2: Defender chooses whether to fight (resist) or concede.

As soon as one party concedes, the other party gets to enjoy the entire resource forever afterwards; thus conflict is less flexible than negotiation.

[9] While T might often be stochastic, we do not see any gain in insight from assuming so.

Types and payoffs Each player can be of two types—*tough* or *normal*. Challenger (Defender) is tough with strictly non-zero probability $\hat{\mu}$ ($\hat{\pi}$). In each period in which Challenger attacks, normal Challenger pays a cost at the rate $\tilde{A} > 0$ and normal Defender suffers a loss at the rate $\tilde{L} > 0$. Tough types have no cost of fighting and continue to do so until the end of time.

We wish to disentangle Defender's desire not to to reveal his type from his fear of offering Challenger an opportunity to intimidate him into conceding. Therefore, we divide the analysis in Section 4 in two cases:

Informed Defender Defender privately observes his type at the beginning of the negotiation phase (at $t = -T$);

Uninformed Defender Defender privately observes his type only if and when conflict begins (at $t = 0$).

Challenger privately observes her type at time $t = -T$ throughout the analysis.

Challenger and Defender discount time at rate r_C and r_D, respectively. We denote the total discounted value of benefits or costs incurred over the interval of length Δ by the same letter, but without the tilde, i.e. c, d, A, and L are the present discounted values over Δ of $\tilde{c}$, $\tilde{d}$, $\tilde{A}$, and $\tilde{L}$. The discounting between successive periods is given by $\delta^i = e^{-r_i \Delta}$, $i = C, D$.

We make a parametric assumption to restrict attention to interesting cases. The first part of Assumption 1 says that the loss from one attack in the next period exceeds the flow utility of the contested resource in the current period. The second part says that the cost of attacking is strictly less than the discounted value of getting the contested resource forever, starting from the current period. It immediately follows from this that Challenger strictly prefers to attack if she knows that one attack will result in Defender conceding immediately. Note that if Assumption 1 fails, then either normal Challenger never attacks or normal Defender never concedes.

Assumption 1 (A1). $\delta^D L > d$; $A < c\left(1 - \delta^C\right)^{-1}$.

Our solution concept is perfect Bayesian equilibrium (henceforth *equilibrium*).

Some remarks on the modelling choices A remark is in order on the two alternative assumptions about information. The first covers a situation where Defender becomes aware of his ability to persist in fighting only when he is involved in conflict (this could be because Challenger's ability to hurt Defender may be unclear). In any case, Defender is aware of his type when conflict starts, because his strategy in the game of conflict depends on his type. In particular this sheds some light on how the likelihood and size of offers vary with the information possessed by parties.

Note that conflict ends the possibility of negotiation. As noted by Langlois and Langlois (2009) the idea that negotiations are common *during* conflict is at odds with what we know about states at war. For example, Pillar (1983) finds that in only nineteen of one-hundred and forty two interstate wars parties negotiated during conflict and before an armistice. Furthermore, once the solution to our game is presented it will be clear that our predictions are not dramatically altered even if we were to allow the parties to negotiate during successive periods of conflict.

2.1 Strategies and Public Beliefs

Notation. Period t of the conflict game lasts from calendar time $(t-1)\,\Delta$ to $t\Delta$; for convenience the subscript refers to the period rather than calendar time during our discussion of conflict. At each period t of conflict, the complete history is summarized by the vector of public beliefs. In stage 1, the state of the game is the vector (μ_{t-1}, π_{t-1}), where μ_{t-1} is Defender's belief that Challenger is tough, and π_{t-1} is Challenger's beliefs that Defender is tough. In stage 2, the state vector is (μ_t, π_{t-1}), where Defender's belief about Challenger's type has been updated from μ_{t-1} to μ_t in light of Challenger's action at stage 1.

Negotiation. In the negotiation phase, Defender's strategy is an offer in $[0,1]$. Normal Challenger's strategy maps from the current beliefs and the offer a probability of accepting the offer in $[0,1]$.

Conflict. In the conflict phase, a (behavior) strategy for normal Challenger is a sequence of mappings $\sigma_t^C : [0,1]^2 \to [0,1]$, $t \in \mathbb{N}$, where $\sigma_t^C(\mu_{t-1}, \pi_{t-1})$ is the probability that the normal type of Challenger concedes (does not attack) in period t as a function of the public beliefs. A strategy for Defender is a sequence of mappings $\sigma_t^D : [0,1]^2 \to [0,1]$, one for each $t \in \mathbb{N}$, where $\sigma_t^D(\mu_t, \pi_{t-1})$ is the probability with which Defender concedes in period t when the public beliefs are μ_t and π_{t-1}.[10] At this point it is natural to wonder why we

[10]Note that Challenger's action in period t depends on the beliefs at the end of period $t-1$, as is standard. In contrast, Defender observes Challenger's move at t, updates her belief about Challenger's type to μ_t, and only then chooses an action.

do not incorporate the history of offers that were made. Note that these offers must have been turned down, for the game would have ended otherwise. Once the public beliefs with which we start conflict have been pinned down we shall show that the equilibrium is determined independently of the actual offers (except in as much as they influenced the probabilities of acceptance and therefore the public belief about Challenger being tough). Hence the notation above is without loss of generality.

Since tough players never concede, the *average* probabilities of concession by Challenger and Defender respectively are obtained by multiplying the probability of the normal type by the probability that the (respective) normal type concedes:

$$\begin{aligned} \bar{\sigma}_t^C(\mu_{t-1}, \pi_{t-1}) &= (1-\mu_{t-1})\,\sigma_t^C(\mu_{t-1}, \pi_{t-1}); \\ \bar{\sigma}_t^D(\mu_t, \pi_{t-1}) &= (1-\pi_{t-1})\,\sigma_t^D(\mu_t, \pi_{t-1}). \end{aligned} \tag{1}$$

Obviously, $\mu_t = 0$ at any history where Challenger concedes. If Challenger has not conceded until period t, the updated belief μ_t that Challenger is tough is derived by Bayes' rule from μ_{t-1} and σ_t^C:

$$\mu_t = \frac{\mu_{t-1}}{1-\bar{\sigma}_t^C(\mu_{t-1}, \pi_{t-1})}. \tag{2}$$

Similarly, if Defender has refused to concede throughout we have

$$\pi_t = \frac{\pi_{t-1}}{1-\bar{\sigma}_t^D(\mu_t, \pi_{t-1})}, \tag{3}$$

with $\pi_t = 0$ if Defender concedes at t.

3 Equilibrium in the Game of Conflict

We solve the game backwards. Before we can solve for the optimal offer in the next section, we must derive the equilibrium of the continuation conflict game, starting with the priors μ_0 and π_0. We do so under an additional assumption, which rules out a negligible set (of zero Lebesgue measure) of beliefs but gives us uniqueness.

Assumption 2 (A2). *The quantities* $\ln \bar{\pi} / \ln \pi_0$ *and* $\ln \bar{\mu} / \ln \mu_0$ *are not integers.*

We first show that optimal strategies must necessarily be of a threshold form.

3.1 Threshold Beliefs

Clearly, each player concedes if he or she believes the other to be tough with probability 1. Since payoffs are continuous in beliefs, so are the optimal strategies. If Defender believes Challenger to be tough with sufficiently high probability, then Defender will concede immediately, even if he knows that normal Challenger will never attack again. Similarly, if Challenger believes Defender to be tough with sufficiently high probability, then she will concede immediately even if she knows that normal Defender will concede with certainty in response to just one more attack. The following lemma finds these thresholds exactly.

Lemma 1 (Threshold beliefs)**.** *Let $\bar{\mu}$ and $\bar{\pi}$ be given by*

$$\bar{\mu} := \frac{1}{\delta^D}\left[1+\left(1-\delta^D\right)\frac{L}{d}\right]^{-1}, \text{ and} \tag{4}$$

$$\bar{\pi} := 1-\left(1-\delta^C\right)\frac{A}{c}. \tag{5}$$

In any equilibrium,

(i) *assuming that normal Defender concedes after a single attack, normal Challenger strictly prefers to attack if $\pi_t < \bar{\pi}$, is just indifferent at $\bar{\pi}$, and strictly prefers not to attack otherwise;*

(ii) *assuming that normal Challenger concedes when Defender does not concede immediately, then Defender strictly prefers not to concede if $\mu_t < \bar{\mu}$, is just indifferent at $\bar{\mu}$, and strictly prefers to concede otherwise.*

Proof. All proofs are in Appendix. □

Later we will use the fact that $\bar{\mu}$ and $\bar{\pi}$ are not functions of priors μ_0 and π_0. These thresholds depend in a very natural way on the discount factors of the players. To illustrate, $\bar{\pi} \to 1$ as Challenger becomes more patient ($\delta^C \to 1$), i.e. Challenger concedes only if Defender is almost certainly tough. Similarly, a very patient Defender concedes only if Challenger is almost certainly tough.[11]

[11] As $\delta^C, \delta^D \to 1$, i.e. as the interval between periods becomes smaller, our model gets closer to the continuous-time models in Ponsatì and Sàkovics (1996) and Abreu and Gul (2000). In the limit, if $\delta^C/\delta^D \to 0$, then (normal) Challenger never attacks; if $\delta^D/\delta^C \to 0$, Challenger attacks with probability 1 in period 1, and (normal) Defender concedes immediately.

3.2 Order of Conflict and the Unique Equilibrium

Depending on parameter values the equilibrium is one of only two mutually exclusive types. Both types have similar qualitative features, analogous to those in continuous-time wars of attrition. First, solely from parameters we can compute a finite number of periods, n, during which normal types of players mimic the corresponding tough type with strictly positive probability. Second, after conflict continues for n periods, normal types concede. Thus conflict goes beyond n periods (indeed, forever) if and only if both players are actually drawn to be the corresponding tough types. We call n the *order* of the conflict.

The order of conflict n naturally relates to the prior beliefs that the players are tough. When both players are believed to be normal with very high probability, then conflict is of high order. Conversely, if at least one player is believed to be tough with high probability, then conflict is of low order. Thus a conflict with low order can be thought of as a conflict between players who are initially believed to be highly committed to fight until ultimate victory.

Definition 1. A conflict is of order $n \in \mathbb{N} \cup \{0\}$ if n is the largest non-negative integer such that $\mu_0 < \bar{\mu}^n$ and $\pi_0 < \bar{\pi}^n$, where the conflict game starts with beliefs (μ_0, π_0) and $\bar{\mu}$ and $\bar{\pi}$ are the thresholds.

Note that the conflict order is 0 if either $\bar{\mu} < \mu_0 < 1$ or $\bar{\pi} < \pi_0 < 1$; otherwise the order is non-zero. In the remaining of the analysis we focus on the case when the conflict order is at least 1; otherwise one of the two players concedes immediately in period 1.

We define the relative commitment of Challenger and Defender as follows; this determines which of the two types of equilibria prevail.

Definition 2. In a conflict of order n, Challenger is *more committed* if $\bar{\mu}^{n+1} < \mu_0 < \bar{\mu}^n$ and $\pi_0 < \bar{\pi}^{n+1}$; Defender is *at least as committed* if $\mu_0 < \bar{\mu}^n$ and $\bar{\pi}^{n+1} < \pi_0 < \bar{\pi}^n$.

Propositions 1 and 2 characterize the unique equilibrium of the conflict phase.

Proposition 1. *If A1 and A2 are satisfied and Challenger is more committed, in the unique equilibrium*

1. *in period* 1*, Challenger attacks with probability* 1 *and Defender concedes with probability*

$$1 - \frac{\pi_0}{\bar{\pi}^n} \frac{1 - \bar{\pi}^n}{1 - \pi_0}; \tag{6}$$

2. *subsequently Challenger and Defender concede with probabilities* $\frac{1-\bar{\mu}}{1-\mu_t}$ *and* $\frac{1-\bar{\pi}}{1-\pi_t}$ *respectively as long as both beliefs are below the respective thresholds, and then concede with probability 1 when the corresponding thresholds are crossed.*

Proposition 2. *If A1 and A2 are satisfied and Defender is at least as committed, in the unique equilibrium is given by*

1. *in period* 1*, Challenger concedes with probability*

$$1 - \frac{\mu_0}{\bar{\mu}^n} \frac{1 - \bar{\mu}^n}{1 - \mu_0}; \tag{7}$$

2. *subsequently Challenger and Defender concede with probabilities* $\frac{1-\bar{\mu}}{1-\mu_t}$ *and* $\frac{1-\bar{\pi}}{1-\pi_t}$ *respectively as long as both beliefs are below the respective thresholds, and then concede with probability 1 when the corresponding thresholds are crossed.*

We give the full proof of propositions 1 and 2 and further details about its intuition in Appendix A.2. Here we offer only the basic thrust behind its logic necessary to understand how negotiation and conflict are linked.

First, from period 1, stage 2 onward, both players concede with positive probability. The fact that normal players are expected to concede with positive probability is what fuels intimidation. Upon observing an attack, Defender must conclude that Challenger is more likely to be tough than he previously believed. Similarly, when Defender does not concede, Challenger must conclude that Defender is more tough than she previously believed.

Second, as each attack from Challenger and each non-concession from Defender increase the public belief that the players are tough, as conflict continues these beliefs get closer and closer to the threshold beliefs $\bar{\mu}$ and $\bar{\pi}$. Once these thresholds are passed, conflict must end.[12] Therefore, the order of conflict n is naturally linked to the maximum number of periods for which conflict protracts.

Third, one key difference between the two cases in propositions 1 and 2 is that when Challenger is more committed, she attacks with probability 1 at the start, whereas she

[12]Lemma 4 in Appendix A.2 shows that says that if both beliefs are strictly below their threshold, no belief leaps over the corresponding threshold at the next step without touching the corresponding threshold exactly. Intuitively, if Defender knows at time t that fighting once will move π_t above its threshold, then he will fight with probability 1. But if he fights with probability 1, then $\pi_t = \pi_{t+1}$. The implication is that when normal types finish conceding at time n, one reputation is exactly at its threshold, i.e. either $\pi_n = \bar{\pi}$ or $\mu_n = \bar{\mu}$.

mixes when Defender is at least as committed as Challenger. That is, when Challenger is perceived to be more likely to be tough, her expected equilibrium payoff from conflict is greater.

Remark 1. If Challenger is more committed than Defender, then normal Challenger's expected payoff is

$$u^C := \left(1 - \frac{\pi_0}{\bar{\pi}^n}\right) \frac{c}{1-\delta^C} - A \tag{8}$$

and normal Defender's expected payoff is $-L$. The unconditional probability of an attack in period $t \in [2, n]$ is given by

$$\Pr\left(\text{attack at } t\right) = \frac{\pi_0}{\bar{\pi}^n} \bar{\mu}^{t-1} \bar{\pi}^{t-2}.$$

Remark 2. If Defender is at least as committed as Challenger, then normal Challenger's expected payoff is 0 and normal Defender's expected payoff is

$$u^D := \left(1 - \frac{\mu_0}{\bar{\mu}^n}\right) \frac{d}{1-\delta^D} - \frac{\mu_0}{\bar{\mu}^n} L. \tag{9}$$

The unconditional probability of an attack in period $t \in [1, n+1]$ is given by

$$\Pr\left(\text{attack at } t\right) = \frac{\mu_0}{\bar{\mu}^n} \left(\bar{\mu}\bar{\pi}\right)^{t-1}.$$

3.3 Vanishingly Small Uncertainty

We explore the limit of our model as uncertainty about players' types becomes small. Our main interest is in determining conditions under which conflict can continue even when both parties are very likely to be normal.

An Example. Consider a symmetric conflict of order n with $\bar{\mu} = \bar{\pi}$ and equal priors. As uncertainty vanishes, the order of conflict grows. The following proposition says that in the limit as $\mu_0 = \pi_0 \to 0$, conflict extends to any period $t \in \mathbb{N}$ with probability at least $\bar{\mu}\left(\bar{\mu}\bar{\pi}\right)^{t-1}$. To give a sense of the magnitude of this lower bound, suppose that a period is a week and the per-week rent from the resource equals $.8$ of the loss from a single attack. For a per-week discount factor $\delta^D = .99$, our result says that we should expect the conflict to extend for an entire year with probability at least $.78$.

Limit Result. From the previous sections, a *conflict* (game) is completely characterized by the prior probabilities of being tough and the thresholds, $(\mu_0, \pi_0, \bar{\mu}, \bar{\pi})$, which is in $(0,1)^4$. Fixing $(\bar{\mu}, \bar{\pi}) \in (0,1)^2$, we can define a sequence of conflict games $\{\Gamma_k\}$ by varying the prior $(\tilde{\mu}_k, \tilde{\pi}_k)$. Let the order of Γ_k be denoted by $n(k)$. We consider sequences that converge to complete information, i.e. $(\tilde{\mu}_k, \tilde{\pi}_k) \to (0,0)$. Obviously, $n(k) \to \infty$ along such sequences.

For such a sequence of conflicts converging towards rationality, the following definition ensures that no player is infinitely more committed (see Definition 2) than the other.

Definition 3. A sequence of conflicts $\{\Gamma_k\}$ is *equipoised* if

$$\lim_{k \to \infty} \frac{\tilde{\mu}_k / \bar{\mu}^{n(k)}}{\tilde{\mu}_k / \bar{\mu}^{n(k)} + \tilde{\pi}_k / \bar{\pi}^{n(k)}} \in (0,1) .$$

Proposition 3 considers a sequence of conflicts such that Defender is at least as committed as Challenger. It says that as uncertainty becomes small, the probability that conflict survives until t is strictly positive for all $t \in \mathbb{N}$.

Proposition 3. *Let $\{\Gamma_k\}$ be an equipoised sequence of conflict games, if either*

1. *Defender is at least as committed for all $k \in \mathbb{N}$, or*
2. *Challenger is more committed for all $k \in \mathbb{N}$,*

then $\liminf_k \ \Pr(\textit{attack at } t) > 0$ *for all* $t \in \mathbb{N}$.

Proof. In Appendix. □

4 Why Negotiation Fails to Eliminate Conflict

We now turn to negotiation and show why in our model conflict might not be avoided even if Defender and Challenger can commit to a peaceful division of the resource. An important feature of our model is that conflict and negotiation are linked: when choosing an action during negotiations, each player knows that her action will affect her opponent's belief that she is tough. In turn, this belief determines the expected payoff from conflict, affecting the relative appeal of negotiating peace.

Recall that the conflict game can have one of two types of equilibria; in this section we focus on the more interesting type where Challenger has a strictly positive utility from attacking given the prior beliefs $\hat{\mu}$ and $\hat{\pi}$, i.e., Challenger is more committed at the prior $\mu_0 = \hat{\mu}$ and $\pi_0 = \hat{\pi}$. We are interested in the case where opportunities to terminate conflict (via concession) arise very frequently; in other words, the gap Δ between successive periods is small. Our main focus is on the probability that conflict begins, i.e., that no peaceful division of the resource is agreed upon by time $t = 0$.

4.1 Uninformed Defender

We begin our analysis of negotiation with the case of an uninformed Defender; this is the appropriate assumption when Defender needs conflict to start before he can assess the real cost of an attack. Thus Defender's offer does not signal his type. In this case, the probability with which conflict begins depends on the discount factors of the two players, relative to the cost parameters of the game. In fact, when Defender is (weakly) more patient than Challenger ($r^D \leq r^C$) negotiations fail and conflict begins with strictly positive probability. This is also the case whenever the interval between negotiations and conflict is sufficiently small.

Proposition 4. *When Defender is uninformed, if*

$$\left(e^{-\left(r^D - r^C\right)T} - 1\right)\frac{r^D}{r^C} < \frac{\hat{\pi}}{1-\hat{\pi}}\frac{1-\hat{\mu}}{\hat{\mu}}\frac{\tilde{A}/\tilde{c}}{\left(\tilde{L}-\tilde{d}\right)/\tilde{d}}, \tag{10}$$

then, for Δ sufficiently small, conflict begins in equilibrium with strictly positive probability even if both Defender and Challenger are normal.

There are two cases of particular interest in which condition (10) holds. First, when Defender and Challenger are equally patient, the left-hand side of (10) is zero, and therefore (10) holds for all T. Second, irrespective of who is more patient, (10) holds whenever T is sufficiently small. Therefore, negotiations can fail if the threat of conflict is imminent.

To explain the intuition behind our main result, suppose that an offer $x \in [0, 1]$ by Defender is accepted by normal Challenger with probability $\alpha(x)$. Then, the total probability of acceptance[13] is $\beta(x) = (1-\hat{\mu})\,\alpha(x)$. An offer is a double-edged sword for Defender. On one hand, a larger offer increases Challenger's incentive to avoid conflict. If this were to be the whole story, then Defender could offer to Challenger exactly Challenger's expected value of conflict and avoid conflict completely. But, on the other hand, there is

[13] We can simply assume that tough Challenger doesn't accept anything short of the whole resource, but we can get this property for free because, as we shall see, Defender will never make an offer that would be large enough to be accepted by tough Challenger.

also a detrimental side to a larger offer: following a rejection, Defender's belief that Challenger is tough increases to $\mu_0 = \hat{\mu}/(1-\beta(x))$, which is increasing in $\beta(x)$. Therefore, offers that Challenger accepts with higher probability raise the belief that Challenger is tough further if rejected, thereby increasing Challenger's expected payoff from conflict.

With this in mind, it is then easy to see why negotiation might fail to avoid conflict. If Defender were to make an offer that normal Challenger would surely accept, a rejection would convince Defender that Challenger was tough for sure. But this implies that the normal Challenger strictly prefers to refuse, as her refusal intimidates Defender into conceding—unless, that is, Defender offers her the discounted value of the whole resource.

For the special case in which Defender discounts time at the same rate as Challenger, we can derive a closed-form expression for the probability with which the optimal offer is accepted in the limiting case of arbitrarily frequent offers ($\Delta \downarrow 0$). Let $\rho^i := e^{-r^i T}$ denote the total discounting from the offer to the start of conflict (from $-T$ to 0). Challenger compares the offer with the payoff of triggering conflict, which itself depends on the offer. If Challenger accepts the offer x with probability $\beta(x)$, she gets $x\tilde{c}/r^C$ from accepting, whereas rejecting x raises the probability of the tough type of Challenger to $\mu_0 = \hat{\mu}/(1-\beta(x))$. Therefore the value of conflict is given by

$$\rho^C P^D(\mu_0, \pi_0)\frac{\tilde{c}}{r^C} = \rho^C P^D\left(\frac{\hat{\mu}}{1-\beta}, \pi_0\right)\frac{\tilde{c}}{r^C}, \tag{11}$$

where the probability $P^D(\mu_0, \pi_0)$ with which Defender concedes as soon as conflict begins is given by the limit as $\Delta \downarrow 0$ of the expression in Proposition 1, Part 1 (proved in Appendix (A.4))

$$P^D(\mu_0, \pi_0) = 1 - \pi_0 \mu_0^{-\sigma^D/\sigma^C} > 0, \text{ where } \sigma^D \equiv \frac{\tilde{A}}{\tilde{c}/r^C}, \ \sigma^C \equiv \frac{\tilde{L}-\tilde{d}}{\tilde{d}/r^D}. \tag{12}$$

Defender's expected discounted share of the resource, as measured from date $-T$, as a function of his offer x and the total acceptance probability β, is

$$V(x) := \beta(1-x) + (1-\beta)\left\{\hat{\pi} + (1-\hat{\pi})\left(1-\rho^D\right)\right\},$$

while his actual discounted utility is $V(x)\frac{\tilde{d}}{r^D}$. The first term is theDefender's portion of the resource after an offer is accepted; the second represents the expected share in case the offer is rejected—the tough type gets the entire share (from time $-T$); the normal type gets the entire resource until the start of conflict and then concedes immediately.

Defender's problem is

$$\max_{x\in[0,1]} \quad V(x) \quad \text{s.t. } x \le \rho^C \left\{ 1 - \hat{\pi} \left(\frac{\hat{\mu}}{1-\beta} \right)^{-(\sigma_D/\sigma_C)} \right\},$$

with the additional restriction that $\beta = 0$ if the inequality is strict. The constraint says that x must be less or equal to Challenger's discounted utility of rejecting an offer that is accepted with probability β, obtained by substituting (12) into (11). Clearly, Defender will never offer a larger share than this. If the constraint is slack, then from Challenger's perspective accepting the offer must be worse than pursuing conflict, and therefore $\beta = 0$.

Let us now focus on the analytically convenient case of equal discounting. Substituting for x in terms of β and letting $\rho^C = \rho^D = \rho$, we obtain

$$\frac{dV(\beta)}{d\beta} = \pi_0 \rho \left\{ \left(\frac{\hat{\mu}}{1-\beta} \right)^{-\sigma^D/\sigma^C} \left(1 - \frac{\beta}{1-\beta} \frac{\sigma^D}{\sigma^C} \right) - 1 \right\}. \tag{13}$$

Equating (13) to zero gives the first-order condition satisfied by the acceptance probability β^* of an optimal offer. As $dV(\beta)/d\beta$ is strictly decreasing in β, the acceptance probability is unique. This proves the following proposition.

Proposition 5. *Let the uninformed Defender discount time at the same rate as Challenger. As* $\Delta \downarrow 0$, *the optimal offer is accepted with probability* $\beta^* \in (0,1)$ *such that*

$$1 - \frac{\beta^*}{1-\beta^*} \frac{\sigma^D}{\sigma^C} = \left(\frac{\hat{\mu}}{1-\beta^*} \right)^{\sigma^D/\sigma^C}. \tag{14}$$

With this in mind, it is now easier to interpret the condition in Proposition 4. First, we note that unequal discounting introduces an additional factor. If negotiation succeeds, it switches consumption from Defender to Challenger. Switching consumption to the less patient player is efficient. Therefore, when Challenger is less patient, the chances for a mutually advantageous agreement are greater. On the contrary, when Challenger is more patient, there is a loss in efficiency in moving consumption to her, and therefore less room for a mutually advantageous agreement.Second, if Defender is sufficiently unlikely to be tough (small $\hat{\pi}$) and Challenger is sufficiently likely to be tough (large $\hat{\mu}$), conflict is so favorable to Challenger and so unfavorable to Defender that Defender prefers to concede the whole discounted value of the resource and conflict is avoided. Finally, when conflict is more costly for Challenger (higher $\tilde{A}$ and smaller $\tilde{L}$), Challenger is relatively more interested in avoiding conflict and it is therefore easier for Defender to make an offer that normal Challenger will accept.

4.2 Informed Defender

We now study the alternative informational assumption where Defender makes the offer with knowledge of his own type. We use the following fact throughout this part: the tough type of Defender must offer 0, since he has no benefit whatsoever from negotiating a settlement. Thus, by making an offer, normal Defender also signals his type. We now show that this signal is strong enough that there is no scope for negotiation whenever Challenger is more patient than Defender. When Challenger is much less patient, it is worthwhile for normal Defender to separate fully from the tough type and resolve the dispute for sure.

Proposition 6. *When Defender is informed, the equilibrium in the limit model is one of three types depending on relative discounting and the probability of the Challenger being tough.*

1. *if and only if $\rho^C \leq \hat{\mu}\rho^D$, then there is a separating equilibrium where conflict between normal types is averted;*

2. *if and only if $\hat{\mu}\rho^D < \rho^C < \rho^D$, then there is a semi-separating equilibrium in which the normal Defender mixes between mimicking the tough type and revealing himself, and conflict between normal types begins with strictly positive probability;*

3. *if and only if $\rho^D \leq \rho^C$, then both types pool on the zero offer and conflict begins with probability 1.*

Notice that separation cannot happen when players are comparably patient (i.e., complete separation occurs only if $\rho^C \leq \hat{\mu}\rho^D < \rho^D$). Therefore, the failure of negotiations under equal discounting is robust to different discount factors.

Conflict can be avoided for sure here when Challenger is sufficiently less patient than Defender. Intuitively, when Challenger is very impatient, she is willing to accept small offers even if she were to know that Defender is normal for sure; in addition negotiation improves efficiency by switching consumption to the more patient player. Therefore, normal Defender has less incentive to conceal his type and makes an offer that normal Challenger accepts with probability 1. But if Challenger is sufficiently patient, normal Defender knows that if he were to reveal his type, Challenger would accept only very large offers. Therefore, normal Defender has an incentive to mimic his tough type and offer $x = 0$. Therefore, conflict begins with positive probability even if both Challenger and Defender are normal. Indeed when Challenger is more patient than Defender, conflict begins for sure. Unless $\rho^C < \hat{\mu}\rho^D$, normal Defender cannot gain from the possibility of negotiations.

A further insight arises from the comparison of the informed and uninformed Defender cases:when ρ^C is slightly larger thanρ^D, Defender does not make an acceptable offer if it is known that he knows his type privately, whereas he makes an acceptable offer when he is known to be uninformed.

5 Brinkmanship (Multiple Offers)

In the previous section, we established why a single round of negotiations fails to avoid conflict. But with multiple rounds, Defender could potentially 'screen' Challenger, i.e., take out a proportion of normal types at each round, lowering the probability of conflict, which could conceivably go to 0 as the number of rounds increases. We now show that this is not in fact the case: our results on the failure of negotiation are robust to an arbitrary number of offers.

There are K times $t_1, \ldots, t_K$, such that $-T = t_1 < \ldots < t_K < 0$, at each of which Defender can offer any fraction of the resource to Challenger. The conflict phase is triggered at time 0 if all K offers are rejected. We focus on the case of equal discounting and uninformed Defender.

Notice that Defender can always afford the same expected payoff he would get if there was only one round of negotiations. In fact, he can choose to make offers that no Challenger would accept until the K-th round, and then make the same offer he would make if there was only one round. Therefore, the question is whether Defender can do better than this by making an acceptable offer before the K-th and last round of negotiation. Proposition 7 says that the answer to this question is negative and that in equilibrium offers are accepted only on the brink of conflict.

Proposition 7. *Let β_k^*, $k = 1, \ldots, K$ be the equilibrium probability that Challenger accepts the offer in round k and β^* be the equilibrium probability that Challenger accepts if there is only one round of negotiation. All equilibria must satisfy*

$$\beta_1^* = \ldots = \beta_{K-1}^* = 0 < \beta_K^* = \beta^*.$$

Once again, the intuition behind our brinkmanship result lies in the ability of Challenger to intimidate Defender—in this case, into making larger offers in the future. In fact, once the last round of negotiation is reached, the continuation game is akin to the game considered in the previous section. By Proposition 4, Point 1, Defender will always make an offer that has positive probability of being accepted. But the size of this offer is increasing in Defender's belief that Challenger is tough. In fact, when Defender believes that Challenger is more likely to be tough, then Challenger's expected payoff from conflict is greater and she accepts only larger offers.

Suppose then that in an earlier round Defender makes an offer that Challenger accepts with positive probability. By rejecting this offer, Challenger raises Defender's belief that she is tough, intimidating her into making a larger offer in Round K. In particular, this offer will have to be larger than the one that Defender would optimally make if he had only one chance to negotiate. In turn, this also means that the earlier offer must be large. In fact, Challenger must be indifferent between accepting the first or the second offer. Thus, the cost for Defender of buying Challenger's agreement is the same in the two periods and it is greater than the cost it pays when there is only one round of negotiation. From Defender's perspective, this is the same as making a single larger offer, which increases both the chances of peace and the cost of it. But since β^* optimally solves the tradeoff between increasing chances of peace and increasing cost of reaching an agreement, then making multiple acceptable offers lowers Defender's utility.

This *brinkmanship* result provides one explanation for negotiation in the shadow of conflict not making no progress until the very last chance. An instructive anecdote comes from the negotiations leading up to the Treaty of Porthsmouth in the Russo-Japanese War (1904–1905); this subsequently gained President Theodore Roosevelt the Peace Nobel Prize. Negotiations for the treaty were held while Russia brought further troops to Manchuria, a move that would have given Russia an advantage in case of conflict. Fredrik (2010) notes how the Japanese delegation demanded the southern part of the island of Sakhalin *and* war reparations throughout the negotiations. Only upon the arrival of further four Russian divisions, at what was conceivably their last chance to do so, did the Japanese drop their claim for reparations and avert conflict.

6 Comparative Statics and Applications

Conflict Duration. We provide some comparative statics regarding the length of conflict, conditional on (i) a first attack, and (ii) a first non-concession by Defender. We first show that the maximum length of a conflict depends on the absolute likelihood that Challenger and Defender are tough. Thus, our model predicts that conflicts between players believed to be tough last for fewer periods.

Corollary 1. *Unless both players are tough, the maximum length of a conflict is determined by the conflict order n. If Challenger is more committed, there is never an attack after period $n+1$. If Defender is at least as committed, there is never an attack after period n.*

Remark. Recall that as uncertainty vanishes, the order of conflict n and its maximum length tend to infinity.

Conditional on the conflict continuing after period 1, the probability that it lasts until $t \leq n$ is independent of which player is more committed or how much the players are likely to be tough. Between period 2 and n, the survival probability of the conflict depends only on the threshold values $\bar{\mu}$ and $\bar{\pi}$, which do not depend on the priors μ_0 and π_0. Thus, until period $t = n$ is reached, conflict is non-duration dependent.

Corollary 2. *In each period $t : 2 \leq t \leq n$, if conflict has not yet ended, Challenger attacks with constant probability $\bar{\mu}$ and Defender does not concede with constant probability $\bar{\pi}$. Therefore, if conflict has not yet ended by the end of period t, the probability that conflict will not end by the end of period $t' : t < t' \leq n$ is $(\bar{\mu}\bar{\pi})^{t'-t}$.*

Thus, after period 1, the probability that conflict ends at time $t > 1$ follows a geometric distribution with Bernoulli probability $1 - \bar{\mu}\bar{\pi}$ until $t = n+1$. We wish to interpret our model as a small deviation from full rationality, i.e. μ_0 and π_0 are small and therefore n is large. As n grows large, the distribution of conflict length beyond period 1 follows a geometric distribution with mean $(1 - \bar{\mu}\bar{\pi})^{-1}$. Therefore, from an empirical perspective, for sufficiently small μ_0 and π_0, conflict is non-duration dependent.

Although the probability that conflict extends to the next period does not depend on time, it nonetheless depend in intuitive ways on other primitives of the model. The following comparative statics follows from the previous corollary and the definition of $\bar{\mu}$ and $\bar{\pi}$. For μ_0 and π_0 sufficiently large such that conflict length is approximated by a geometric distribution, the same comparative statics apply to the expected length of conflict.

Corollary 3. *Conditional on there being a conflict at time $t > 1$, the probability of an attack in period $t' > t$ is increasing in the value of the resource (c and d) and the players patience (δ^C and δ^D); it is decreasing in the cost of fighting (A and L).*

Prediction 1. Conflict is non-duration dependent. Conflict duration is determined by the military capabilities of the parties and the value of the resource at stake: it decreases

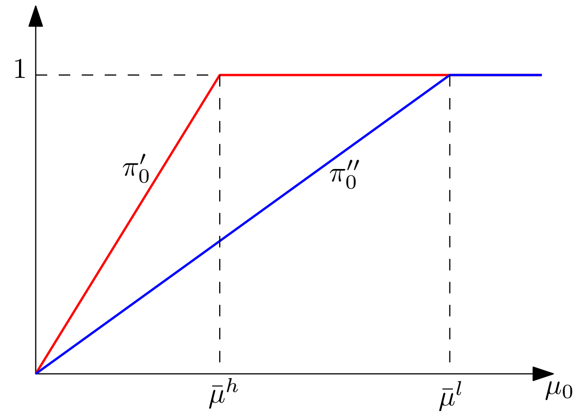

Figure 1: Probability of First Attack

with the cost of fighting and increases with the value of the resource. Only if the parties are tough with sufficiently large probability, then conflict duration is sensitive to these probabilities: it decreases with the probability that at least one party is tough.

Probability of Conflict. We now turn to the question of when a conflict is more likely to begin, i.e. there is a first attack. Figure 1 depicts the total probability of first attack as a function of μ_0 for two different values of π_0, $\pi_0' \in (\bar{\pi}^{h+1}, \bar{\pi}^h)$ and $\pi_0'' \in (\bar{\pi}^{l+1}, \bar{\pi}^l)$. In both cases, the probability of first attack is strictly increasing for low values of μ_0 and equals 1 for higher values.

Corollary 4. *Fix the likelihood π_0 that Defender is tough. The probability that Challenger begins to attack is increasing in the likelihood of Challenger being tough μ_0. It is* strictly *increasing if and only if Defender is at least as committed as Challenger.*

For Defender, an image of toughness can pay: if Defender is at least as committed as Challenger, then the probability of a conflict is strictly less than 1. In this case, the probability of a conflict is $\mu_0/\bar{\mu}^n$, where n is the largest natural number such that $\pi_0 \leq \bar{\pi}^n$. Thus, if π_0 increases, the probability of conflict decreases. Figure 1 depicts the total probability of first attack as a function of μ_0 for $n = h$ (blue line) and $n = l < h$ (red line). Increasing π_0 from π_0' to $\pi_0'' > \pi_0'$ moves the line representing the probability of first attack to the right.

Corollary 5. *Let Defender be at least as committed as Challenger. Then, the probability that Challenger begins to attack is decreasing in π_0.*

The advantage of being perceived as tough should not be overstated. Recall that after the first attack, the expected payoff for Defender is $-L$, independently of π_0. Indeed, in equilibrium, Defender is indifferent between conceding and resisting whenever he plays.

Prediction 2. The probability that a conflict begins depends on the relative probability that the parties are tough. In particular, if Challenger is more committed than Defender, then a greater difference between the probabilities of being tough increases the probability that conflict begins.

6.1 Democracy and War

One of the few empirical regularities about conflict is that pairs of democracies are less likely to fight each other.[14] Prediction 2 suggests that conflict begins when there is greater unbalance between the parties' probability of being tough. We argue that democratic leaders are kept in check by their citizens and that therefore democracies tend to have similar probabilities of being tough or irrational. On the contrary, autocrats of the like of Kim Jung-un are known for their unpredictable behavior. Thus, a pair of autocracies or one democracy and one autocracy are more likely to have unbalanced probabilities of being tough and are therefore more likely to engage in conflict.

The evidence concerning the effect of democracy on the duration of conflict is ambiguous. Bennet and Stam (2009), Langlois and Langlois (2009), and Henderson and Bayer (2013) find that the relation between democracy and conflict duration is not significant once the military capabilities of the parties and the physical characteristics of conflict (common boundaries, terrain, etc.) are taken into account. Prediction 1 suggests that that the probability of continuation of conflict depends indeed only on the costs and benefits of war, and only to a lesser extent on the probability of being tough. Thus, in our model as in the real world, physical and technological characteristics matter more than political ones.

6.2 Terrorism

Terrorist conflicts can last from a few months to many decades. A common distinction in the literature is the one between religious and secular terrorism. Religious terrorist groups are believed to be fiercely committed to an ultimate goal to the exclusion of compromise. For example, Taheri (1987) notes that a key difference between Islamic and secular terrorism is that the first "is clearly conceived and committed as a form of Holy War which

[14] See Bueno de Mesquita and Smith (2012) for a survey.

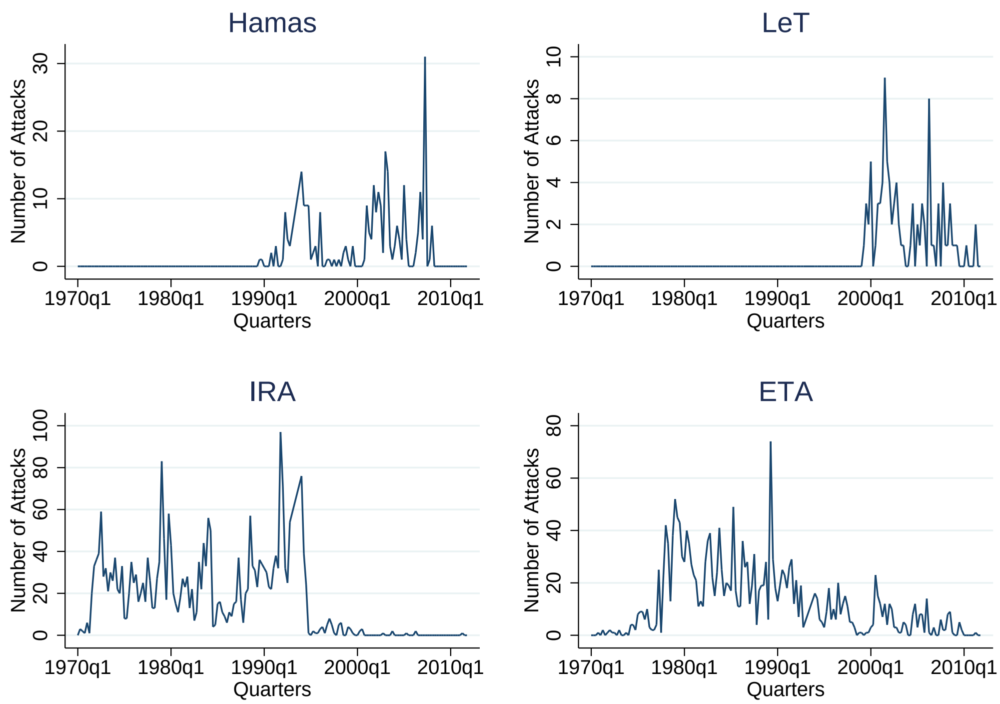

Figure 2: Number of Attacks by Quarter of Major Terrorist Groups

can only end when total victory has been achieved" (p. 7) (see also Hoffman, 1995). With this interpretation, predictions 1 and 2 say that religious terrorist groups are more likely to start conflict, but that such conflicts will be of shorter duration when they start.

Figure 2 depicts the number of attacks by year-quarter for four major terrorist groups from 1970 to 2012.[15] All four terrorist groups are ultimately motivated by a separatist goal. Hamas and Lashkar-e-Taiba (LeT) are Islamic terrorist groups originating respectively in Palestine and Afghanistan (active in Pakistan and India). The (Provisional) Irish Republican Army (IRA) and Basque Fatherland and Freedom (ETA) are separatist military groups respectively in Ireland and the Basque Country in Spain and France. The pattern of conflicts is consitent with our model. Both religious groups, Hamas and LeT, show short burst of conflicts followed by periods of little or no activity. The time-series of Hamas's attacks, in particular, shows clearly three brief periods of conflict corresponding to the First Intifada (from Hamas foundation in 1988 to 1993), the Second Intifada (September 2001 to 2005) and the conflict culminated with the Gaza War (June 2006 to January

[15] Our calculations on data from National Consortium for the Study of Terrorism and Responses to Terrorism (START), 2012, Global Terrorism Database [globalterrorismdb_1012dist]. Retrieved from http://www.start.umd.edu/gtd.

2009). In stark contrast, both secular groups, the IRA and ETA, show no such a pattern. The Troubles in Ireland and the UK which begun in 1969 lasted almost uninterrupted until the late 90s, when the Good Friday Agreements were signed (April 1998). Similarly, ETA's attacks continued uninterrupted from the assassination of of Admiral Luis Carrero Blanco (Franco's chosen successor) in 1973 to the *Spirit of Ermua* massive demonstrations against ETA in 1997. A second conflict started with a car bomb on January 21, 2000.

6.3 Workers' Unions and Macroeconomic Conditions

When workers go on strikes, they effectively incur a cost for themselves *and* impose a cost on their employers. The implicit threat to the employers is that the working conditions are so poor that the workers would lose their current salary rather than continuing to work. The committed type workers are therefore those who effectively have a sufficiently good outside option—typically, another employment. On the other hand, employers deny a concession to signal that they are unable to raise pay without jeopardizing the future of their firms. A committed employer is thus one who prefers losing human capital to raising pay.

The relevant stylized empirical facts are that strikes (i) are more likely in good economic times, and (ii) longer in bad economic times (Kennan and Wilson, 1990). Both facts agree with the predictions of our model.

In a typical business cycle, in good times, workers are more likely to have a good outside option—μ_0 is larger—and businesses are more likely to be able to raise pay—π_0 is smaller. Thus, our model predicts that strikes are more likely in good economic times (Challenger is relatively more likely to be committed) than in bad times. Furthermore, the cost of a strike is greater (at least for the businesses) in good economic times, i.e. L is greater in good economic times. Thus, our model predicts longer strikes in bad times than in good ones.

7 Conclusion

We provide a framework to make sensible predictions about 'wars of intimidation', where a challenger repeatedly attacks to convince the defender that more attacks will come if the latter doesn't concede. The model is built on slight probabilities that each player could be a type that faces no cost of fighting, and persists with conflict until the opponent concedes. Our model predicts prolonged conflicts even when uncertainty about the players' type is arbitrarily small. Our comparative static predictions find support in applications of interest in economics and political science.

Wars of intimidation begin when negotiations fail. Our model allows us to uncover a novel reason for why this happens: offers that have a higher probability of being accepted also increase the incentive for the aggressor to initiate conflict. Thus, negotiations can mitigate but not prevent conflict. This result should not be seen as a claim that negotiations are useless: a neutral observer who seeks to reduce conflict will gain from bringing the parties to the negotiation table. If negotiations succeed, then conflict is avoided; if they fail, conflict will be shorter. On the other hand, our brinkmanship result shows that there is little point in insisting on multiple rounds of negotiations. If acceptable offers are to be made, they will only be made at the last opportunity to avert conflict. In a sense, neutral observers should take advantage of any ultimatum imposed by the parties, rather than pressuring them to have softer deadlines.

References

[1] **Abreu, Dilip and Faruk Gul (2000).** Bargaining and Reputation, *Econometrica,* 68(1): 85–117.

[2] **Baliga, Sandeep and Tomas Sjöström (2013).** Bargaining and War: A Review of Some Formal Models, *The Korean Economic Review*, 29(2): 235–266.

[3] **Bennet, Scott and Allan C. Stam (2009).** Revisiting Predictions of War Duration, *Conflict Management and Peace Science*, 26(3): 256–267.

[4] **Brito, Dagobert L. and Michael D. Intriligator (1985).** Conflict, War and Redistribution, *The American Political Science Review*, 79(4): 943–957.

[5] **Bueno de Mesquita, Bruce and Alastair Smith (2012).** Domestic Explanations of International Relations, *Annual Review of Political Science*, 14: 1–21.

[6] **Chatterjee Kalyan and William Samuelson (1983).** Bargaining under Incomplete Information, *Operations Research*, 31(5): 835–851

[7] **Chatterjee Kalyan and Larry Samuelson (1988).** Bargaining under two-sided incomplete information: The unrestricted offers case. *Operations Research*, 36(4): 605–618.

[8] **Cho, In-Koo (1990).** Uncertainty and Delay in Bargaining, *Review of Economic Studies*, 57: 575–596.

[9] **Crawford, Vincent P. (1982).** A Theory of Disagreement in Bargaining, *Econometrica*, 50(3): 607–637.

[10] **Esteban, Joan and József Sákovics (2008).** A Theory of Agreements in the Shadow of Conflict: The Genesis of Bargaining Power, *Theory and Decision*, 65(3): 227–252.

[11] **Fearon, James D. (1995).** Rationalist Explanations for War, *International Organization*, 49(3): 379–414.

[12] **Fearon, James D. (2013).** Fighting Rather Than Bargaining, *mimeo*, Stanford University.

[13] **Fudenberg, Drew, David Levine and Jean Tirole (1985).** Infinite horizon models of bargaining with one-sided incomplete information. In: Roth, Al (ed.) *Game theoretic models of bargaining*. Cambridge: Cambridge University Press.

[14] **Gul, Faruk, Hugo Sonnenschein and Robert Wilson (1986).** Foundations of Dynamic Monopoly and the Coase Conjecture, *Journal of Economic Theory* 39: 155–190.

[15] **Heifetz, Aviad and Ella Segev (2005).** Escalation and Delay in Protracted International Conflicts, *Mathematical Social Sciences*, 49: 17–37.

[16] **Henderson, Errol A. and Reşat Bayer (2013).** Wallets, Ballots, or Bullets: Does Wealth, Democracy, or Military Capabilities Determine War Outcomes?, *International Studies Quarterly*, 57(2): 303–317.

[17] **Hendricks, Ken, Andrew Weiss and Charles Wilson (1988).** The War of Attrition in Continuous Time with Complete Information. *International Economic Review*, 29(4), 663–680.

[18] **Hodler, Roland and Dominic Rohner (2012).** Electoral Terms and Terrorism, *Public Choice*, 150(1): 181–193.

[19] **Hoffman, Bruce (1995).** "Holy Terror": The Implications of Terrorism Motivated by Religious Imperatives, *Studies in Conflict & Terrorism*, 18(4): 271–284.

[20] **Hörner, Johannes and Nicolas Sahuguet (2010).** A War of Attrition with Endogenous Effort Levels, *Economic Theory*, 47: 1–27.

[21] **Kennan, John and Robert Wilson (1990).** Can Strategic Bargaining Models Explain Collective Bargaining Data?, *The American Economic Review*, 80(2): 405–409.

[22] **Kornhauser, Lewis, Ariel Rubinstein and Wilson, Charles (1989).** Reputation and Patience in the'War of Attrition'. *Economica*, 15–24.

[23] **Kreps, David and Robert Wilson (1982).** Reputation and imperfect information. *Journal of Economic Theory*, 27, 253–79.

[24] **Langlois, Catherine C. and Jean-Pierre P. Langlois (2009).** Does Attrition Behavior Help Explain the Duration of Interstate Wars? A Game Theoretic and Empirical Analysis, *International Studies Quarterly*, 53(4): 1051–1073.

[25] **Lapan, Harvey E. and Todd Sandler (1988).** To bargain or not to bargain: That is the question, *The American Economic Review*, 78(2): 16–21.

[26] **Maynard-Smith, John (1974).** The Theory of Games and the Evolution of Animal Conflicts, *Journal of Theoretial Biology*, 47: 209–22.

[27] **Milgrom, Paul and John Roberts (1982).** Predation, Reputation, and Entry Deterrence, *Journal of Economic Theory*, 27, 280–312.

[28] **Ordover, Janusz A. and Ariel Rubinstein (1986).** A Sequential Concession Game with Asymmetric Information, *The Quarterly Journal of Economics*, 101(4): 879–888.

[29] **Pillar, Paul R. (1983).** *Negotiating Peace: War Termination as a Bargaining Process*, Princeton: Princeton University Press.

[30] **Ponsatì Clara and József Sákovics (1995).** The war of attrition with incomplete information, *Mathematical Social Sciences*, 29(3), pages 239-254.

[31] **Powell, Robert (1996).** Bargaining in the Shadow of Power, *Games and Economic Behaviour*, 15: 255–289.

[32] **Powell, Robert (2004).** Bargaining and Learning While Fighting, *American Political Science Review*, 48(2):344–361.

[33] **Reiter, Dan (2003).** Exploring the Bargaining Model of War, *Perspectives on Politics*, 1(1):27–43.

[34] **Sanchez-Pages, Santiago (2009a).** Conflict as a part of the bargaining process, *Economic Journal*, 119: 1189–1207.

[35] **Sanchez-Pages, Santiago (2009b).** Bargaining and Conflict with Incomplete Information, *working paper*.

[36] **Schelling, Thomas C. (1956).** An Essay on Bargaining, *American Economic Review*, 6: 281–306.

[37] **Schelling, Thomas C. (1960).** *The Strategy of Conflict*, Cambridge, MA: Harvard University Press.

[38] **Schelling, Thomas C. (1966).** *Arms and Influence*, New Haven, CT: Yale University Press.

[39] **Silverman, Dan (2004).** Street Crime and Street Culture, *International Economic Review*, 45(3): 761–786.

[40] **Sobel, Joel and Ichiro Takahashi (1983).** A Multistage Model of Bargaining, *Review of Economic Studies*, 50: 411–426.

[41] **Stanton, Fredrik (2010).** *Great Negotiations: Agreements that Changed the Modern World*, Yardley, Pennsylvania: Westholme Publishing.

[42] **Taheri, Amir (1987).** *Holy Terror: The Inside Story of Islamic Terrorism*, Sphere Books, London.

[43] **Waltz, Kenneth N. (1979).** *Theory of International Politics*, Reading, MA: Addison-Wesley.

[44] **Yared, Pierre (2009).** A Dynamic Theory of War and Peace, *Journal of Economic Theory*, 145: 1921–1950.

A Omitted Proofs

A.1 Proof of Lemma 1

Proof. Challenger's payoff from choosing $\sigma^C_{t+1}(\mu_{t+1}, \pi_t) = 0$, i.e. attacking for sure at $t+1$, is maximum if after an attack at $t+1$ the normal type of Defender concedes for sure, *and* if Challenger concedes at $t+2$ if Defender does not concede at $t+1$. Note that when Defender concedes in period $t+1$, Challenger gets a flow payoff of c from the resource starting with the current period. When Defender is the commitment type, he doesn't concede at $t+1$ and Challenger gets none of the resource because she concedes at $t+2$. Hence Challenger's *maximum* payoff from attacking at $t+1$ is

$$-(1-\delta^C)A + (1-\pi_t)c + \pi_t \cdot 0. \tag{15}$$

Expression (15) is zero at $\bar{\pi}$, and negative above it. This proves part (i).

For part (ii) note that $\sigma^D_{t+1}(\mu_t, \pi_t) = 1$ implies that if Defender does not concede in period $t+1$ then Challenger must put probability 1 on the tough type, i.e. $\pi_t = 1$ and hence will find it optimal to concede at $t+2$ by step (i) above. Then the payoff of attacking is given by (15) and the payoff of conceding is 0; hence $\bar{\pi}$ is again the point of indifference.

Part (iii) follows from a similar argument to part (i). □

A.2 Equilibrium in the Game of Conflict (Propositions 1 and 2)

We now present a few lemmas that identify necessary conditions for equilibrium and thereby pin down the unique one in the game of conflict.

We first ask if a normal player ever mimics the tough type. Lemma 2 says that only Challenger in period 1 can mimic the tough type (i.e. attack) with probability 1.

Lemma 2. *In any equilibrium, normal types of both players concede with strictly positive probability in all periods, except possibly Challenger in period* 1*.*

Proof. The *concession sequence* $\langle \kappa_i \rangle_{i \in \mathbb{N}}$ of any strategy profile is a sequence in $[0,1]$, where each odd (even) term is the probability that Challenger (respectively, Defender) concedes at that time conditional on no player having conceded yet. A concession sequence arising from an equilibrium profile is called an *equilibrium concession sequence*.

Lemma 2 then says that in any equilibrium concession sequence, all terms (except possibly the first) must be strictly positive.

Step 1. The proof is based on the key idea that if the string $(\kappa_i, 0, \kappa_{i+2})$ appears in an equilibrium concession sequence and $\kappa_{i+2} > 0$, then $\kappa_i = 1$: if the opponent is not conceding in the interim the value of concession can only go down because the positive cost to fighting strictly exceeds the flow utility derived from the resource; therefore concession should have been strictly better at the step before.

Step 2. We now show that, along any concession sequence, adjacent terms cannot be 0. Let $\kappa_i = 0 = \kappa_{i+1}$; if $\kappa_{i+2} > 0$, it would contradict Step 1. Induction implies that if two adjacent terms of the concession sequence are 0, all subsequent terms are 0 too. But since there is a positive probability of the tough type, it cannot be an equilibrium to never concede, knowing that your opponent will not. Therefore, no equilibrium concession sequence contains adjacent 0s.

Step 3. Suppose $\kappa_i = 0$ for some $i > 1$. By Step 2, we must have $\kappa_{i+1} > 0$; from Step 1 it means that $\kappa_{i-1} = 1$. If the player who is supposed to concede with probability 1 does not do so, his/her reputation immediately jumps to 1 and the normal opponent must concede immediately thereafter, i.e. $\kappa_i = 1$—a contradiction! □

Lemma 3 characterizes the players' strategies if they are indifferent for two consecutive periods.

Lemma 3. *If Challenger is indifferent between conceding at times t and $t+1$ in any equilibrium, then Defender's equilibrium concession probability and the public beliefs about him are*

$$\tilde{\sigma}^D(\pi_{t-1}) := \frac{1-\bar{\pi}}{1-\pi_{t-1}}; \text{ and } \pi_t = \frac{\pi_{t-1}}{\bar{\pi}} \tag{16}$$

respectively. Similarly, if Defender is indifferent between conceding at times t and $t+1$ in any equilibrium, then Challenger's probability of conceding and the public beliefs about her type are

$$\tilde{\sigma}^C(\mu_t) := \frac{1-\bar{\mu}}{1-\mu_t}, \text{ and } \mu_{t+1} = \frac{\mu_t}{\bar{\mu}}. \tag{17}$$

Proof. Challenger is indifferent if and only if $V_t^C(\mu_{t-1}, \pi_{t-1}) = 0$, which is the payoff of Challenger from conceding. Suppose that Challenger is indifferent for two consecutive periods: $V_t^C(\mu_{t-1}, \pi_{t-1}) = V_{t+1}^C(\mu_t, \pi_t) = 0$. By Lemma 2, $\sigma_{t+1}^C(\mu_t, \pi_t) \neq 1$. Therefore

$$\bar{\sigma}_t^D(\mu_t, \pi_{t-1}) = 1 - \bar{\pi}.$$

This corresponds to a strategy for Defender such that $\sigma_t^D(\mu_t, \pi_{t-1}) = \tilde{\sigma}^D(\pi_{t-1})$ in (16) if $\pi_{t-1} \leq \bar{\pi}$.

Defender is indifferent if and only if $V_t^D(\mu_t, \pi_{t-1}) = -L$, the payoff Defender gets if he concedes.[16] Suppose that Defender is indifferent for two consecutive periods (or is indifferent at time t and concedes at time $t+1$): $V_{t+1}^D(\mu_{t+1}, \pi_t) = -L = V_t^D(\mu_t, \pi_{t-1})$. By Lemma 2, $\sigma_t^D(\mu_t, \pi_{t-1}) \neq 1$. Therefore

$$\bar{\sigma}_{t+1}^C(\mu_t, \pi_t) = 1 - \bar{\mu}.$$

This corresponds to a strategy for Challenger such that $\sigma_{t+1}^C(\mu_t, \pi_t) = \tilde{\sigma}^C(\mu_t)$ in (17) if $\mu_t \leq \bar{\mu}$. □

Remark 3. Challenger's mixing probability at $t = 1$ need not equal $\tilde{\sigma}^C$; Defender's mixing at $t = 1$ can be different from $\tilde{\sigma}^D$ only if Challenger strictly prefers to attack at $t = 1$.

[16]This payoff is not 0 but $-L$ for the same reason that the term $-L$ appears in equation (**??**)—when it is Defender's turn to decide if he wants to concede or prolong the fight, Challenger has already attacked and the loss will be experienced by Defender in the current period regardless of his choice of move.

Combining Lemmas 2 and 3, in equilibrium both players concede with the probabilities in Lemma 3 above, except possibly at $t = 1$. Beliefs evolve according to the above lemma, except possibly at $t = 1$ and until they hit $\bar{\mu}$ or $\bar{\pi}$.

Conflict continues as long as no player has conceded. If players mix, then beliefs about their type increase until a threshold is crossed.

Lemma 4 says that if both beliefs are strictly below their threshold, no belief leaps over the corresponding threshold at the next step without touching the corresponding threshold exactly.

Lemma 4. *In equilibrium (i)* $\pi_t < \bar{\pi}$ *and* $\mu_{t+1} < \bar{\mu}$ *implies* $\pi_{t+1} \leq \bar{\pi}$ *; (ii)* $\mu_t < \bar{\mu}$ *and* $\pi_t < \bar{\pi}$ *implies* $\mu_{t+1} \leq \bar{\mu}$.

Proof. Suppose not. Let $\pi_t < \bar{\pi}$, $\mu_{t+1} < \bar{\mu}$ but $\pi_{t+1} > \bar{\pi}$. Lemma 1 implies that normal Challenger will concede with probability 1 at time $t+2$ if Defender does not concede at $t+1$. So if Defender does not concede at time $t+1$ he gets a continuation payoff of d from $t+2$ onwards if Challenger is the normal type; since Challenger is normal with probability $1-\mu_{t+1}$, Defender's payoff from $t+1$ (the current period) onwards is

$$\left(1-\delta^D\right)(d-L)+\delta^D\left[\left(1-\mu_{t+1}\right)d+\mu_{t+1}\left(-L\left(1-\delta^D\right)+0\cdot\delta^D\right)\right].$$

Defender strictly prefers to not concede if the above exceeds the payoff $-\left(1-\delta^D\right)L$ from conceding immediately at $t+1$:

$$\left(1-\delta^D\right)d+\delta^D\left[\left(1-\mu_{t+1}\right)d-\mu_{t+1}L\left(1-\delta^D\right)\right]>0. \tag{18}$$

Inequality (18) reduces to $\mu_{t+1} < \bar{\mu}$, which is true by assumption. Therefore Defender strictly prefers to fight at $t+1$, i.e. $\sigma^D_{t+1}\left(\mu_{t+1}, \pi_t\right) = 0$—which contradicts Lemma 2, implying that $\pi_t < \bar{\pi}$ and $\mu_{t+1} < \bar{\mu}$ cannot lead to $\pi_{t+1} > \bar{\pi}$.

Let $\mu_t < \bar{\mu}$ and $\pi_t < \bar{\pi}$, but $\mu_{t+1} > \bar{\mu}$. By a similar logic Challenger strictly prefers to fight at $t+1$ if

$$-\left(1-\delta^C\right)A+\left(1-\pi_t\right)c+\pi_t\cdot 0>0.$$

The net utility for Challenger to fight at period t is $-A$ and with probability $1-\pi_t$ Defender will concede and Challenger will get c forever. The expression above reduces to $\pi_t < \bar{\pi}$. So Challenger strictly prefers to fight at $t+1$, i.e. $\sigma^C_{t+1}\left(\mu_t, \pi_t\right) = 0$—which contradicts Lemma 2. □

What do our previous results imply about period 1's probability of attack? By Lemma

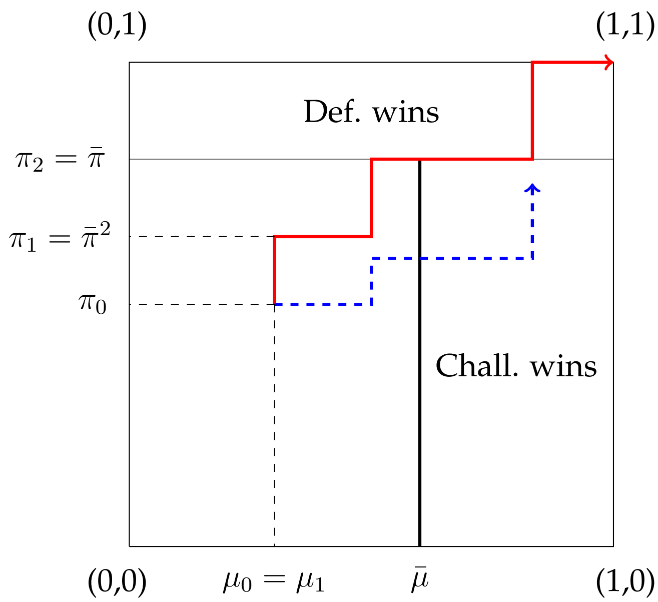

Figure 3: How Beliefs Evolve

3, from period 2 onward beliefs must grow by a factor $\bar{\mu}^{-1}$ and $\bar{\pi}^{-1}$, respectively. The solid line in Figure 3 depicts the equilibrium evolution of beliefs in a conflict of order 2 with Challenger more committed than Defender. The dashed line represents the evolution of beliefs if it is common knowledge that both Defender and Challenger play the strategies in lemma 3 from period 1 onward. In this case, $\pi_1 < \bar{\pi}$ and $\mu_2 > \bar{\mu}$, violating Lemma 4. In equilibrium, Defender must concede with sufficiently large probability in period 1 so as to 'level the playing field' with Challenger and guarantee $\pi_2 = \bar{\pi}$. Since Defender is conceding with a higher probability than what would make Challenger indifferent, in period 1 Challenger strictly prefers to attack.

Remark 4. If Challenger is more committed than Defender, in period 1 Challenger attacks with probability 1 and Defender concedes with probability $1 - \pi_0/\bar{\pi}^n > 1 - \bar{\pi}$.

A similar logic applies to the case when Defender is at least as committed as Challenger. In this case, Challenger must concede with sufficiently high probability in period 1 so as to 'level the playing field' with Defender and guarantee $\mu_n = \bar{\mu}$.

Remark 5. If Defender is at least as committed as Challenger, in period 1 Challenger attacks with probability $\mu_0/\bar{\mu}^n$.

The next lemma shows that along the equilibrium path, provided no one concedes, both reputations grow according to equations (2) and (3) from period 2 onwards until a time t when either $\mu_t = \bar{\mu}$ or $\pi_t = \bar{\pi}$.

Lemma 5. *For any period $t \geq 2$, if $\pi_{t-1} \leq \bar{\pi}$ and $\mu_t \leq \bar{\mu}$, then Challenger plays $\tilde{\sigma}^C(\mu_t)$ and Defender plays $\tilde{\sigma}^D(\pi_{t-1})$.*

Proof. We show the result for Defender. The result for Challenger follows a symmetric argument.

Proceed by contradiction. If $\sigma_t^D(\pi_t) \neq \tilde{\sigma}^D(\pi_t)$, by Lemma 3, Challenger is not indifferent at either t or at $t+1$. There are two possibilities. First, she strictly prefers to concede. But then Defender would concede with probability 0 in the previous period, contradicting Lemma 2. Second, she strictly prefers to fight. But then by Lemma 2 she is Challenger in period 0 and $t = 1 < 2$. □

The lemma above is useful in proving Propositions 1 and 2, which apply, respectively, to the cases where Challenger is more committed than Defender and where Defender is at least as committed as Challenger.

A.2.1 Proof of Proposition 1

Proof. Existence. We first show that the stategies σ^* defined in Proposition 1 constitute an equilibrium. From Lemma 3 it is clear that after the first move by Challenger in period 1 players are indifferent and therefore willing to mix. Since normal players concede in σ^* once the thresholds are crossed, this is consistent with Lemmas 1. Since Defender concedes with a *larger* probability than $\tilde{\sigma}^D$ in the first period, Lemma 3 implies that Challenger strictly prefers to fight at $t = 1$. Also note that by Bayes' rule the equilibrium belief about Challenger's type after non-concession at $t = 1$ is given by $\bar{\mu}^n$.

Uniqueness. If $\mu_0 \geq \bar{\mu}$, then Lemma 1 implies that the above is the only equilibrium; similarly for the case $\pi_0 \geq \bar{\pi}$. Therefore let $(\mu_0, \pi_0) < (\bar{\mu}, \bar{\pi})$, so that $n \geq 1$. If normal types follow $\tilde{\sigma}^C, \tilde{\sigma}^D$ defined in equations (17) and (16) up to and including time n, there will be a jump since $\pi_0/\bar{\pi}^n > \bar{\pi}$; but jumps are ruled out by Lemma 4. By Lemmas 3 and 5, the only freedom we have is in choosing different strategies for $t = 1$.

By contradiction, suppose that Challenger concedes with positive probability in period 1. This implies she expects Defender to concede with probability at least $\tilde{\sigma}^D$. But this implies that there is $m \leq n$ such that beliefs are (μ_{m+1}, π_m) with $\mu_{m+1} > \bar{\mu}$ and $\pi_m < \bar{\pi}$, contradicting Lemma 4.

Last, since Challenger cannot concede with probability less than 0, we have that $\mu_{n+1} > \bar{\mu}$. Thus, by Lemma 4, Defender must concede in period 1 with probability exactly σ_t^{D*}. □

A.2.2 Proof of Proposition 2

Proof. Existence. As before, it can be checked that Lemmas 1 and 3 imply that the above is an equilibrium. In particular, σ_1^{C*} and Bayes' rule imply that the equilibrium belief about Challenger's type after non-concession at $t = 1$ is given by $\bar{\mu}^n$.

Uniqueness. If $\mu_0 \geq \bar{\mu}$, then Lemma 1 implies that the above is the only equilibrium; similarly for the case $\pi_0 \geq \bar{\pi}$. Therefore let $(\mu_0, \pi_0) < (\bar{\mu}, \bar{\pi})$, so that $n \geq 1$. If normal types follow $\tilde{\sigma}^C, \tilde{\sigma}^D$ defined in equations (17) and (16) up to and including time n, there will be a jump since $\pi_0/\bar{\pi}^n > \bar{\pi}$; but jumps are ruled out by Lemma 4. By Lemmas 3 and 5, the only freedom we have is in choosing different strategies for $t = 1$.

Case 1: $\sigma_1^C < \sigma_1^{C*}$. Suppose that $\sigma_1^C < \sigma_1^{C*}$. The inequality $\sigma_1^C < \sigma_1^{C*}$ implies that Challenger's reputation increases at a slower rate such that $\mu_n < \bar{\mu}$. If $\sigma_1^D < \tilde{\sigma}_1^D$, then Challenger prefers to concede immediately ($\sigma_1^C = 1$) since Challenger is just indifferent at $\tilde{\sigma}^D$; this contradiction implies that $\sigma_1^D \geq \tilde{\sigma}^D$, which in turn gives $\pi_1 \geq \pi_0/\bar{\pi}$ and therefore $\pi_n > \bar{\pi}$ i.e. there exists $m \leq n$ such that belief profile is (μ_m, π_m) with $\mu_m < \bar{\mu}$ and $\pi_m > \bar{\pi}$, contradicting Lemma 4. Therefore, $\sigma_1^C \geq \sigma_1^{C*}$ is the only possibility in equilibrium.

Case 2: $\sigma_1^C > \sigma_1^{C*}$. Suppose that $\sigma_1^C > \sigma_1^{C*}$. Now $\mu_1 > \mu_0/\bar{\mu}$, $\mu_2 > \mu_0/\bar{\mu}^2$, etc. Since Proposition 1 implies that Defender's reputation is growing as the same rate $1/\bar{\pi}$ it follows from the Definition 2 and $\bar{\mu}^{n+1} < \mu_0$ that $\mu_n > \bar{\mu}$, i.e. a jump occurs by time n. Therefore, $\sigma_1^C \leq \sigma_1^{C*}$ is the only possibility in equilibrium.

Last, since Challenger must be indifferent at $t = 0$ to play σ_1^{C*}, then $\sigma_1^D = \tilde{\sigma}^D = \sigma_1^{D*}$. ☐

A.3 Proof of Proposition 3

Proof. Part 1. Defender is at least as committed as Challenger is: $\frac{\tilde{\pi}_k}{\bar{\pi}^{n(k)}} \in (\bar{\pi}, 1)$. Thus, $\liminf \frac{\tilde{\pi}_k}{\bar{\pi}^{n(k)}} \geq \bar{\pi}$. By Proposition 2, for any $k \in \mathbb{N}$, the equilibrium probability of attack at time $t \in \mathbb{N}$ is given by

$$\Pr(\text{attack at } t) = (\bar{\mu}\bar{\pi})^{t-1} \frac{\tilde{\mu}_k}{\bar{\mu}^{n(k)}}, \text{ for all } t \in [1, n(k)].$$

$$\begin{aligned} \therefore \liminf \Pr(\text{attack at } t) &= (\bar{\mu}\bar{\pi})^{t-1} \liminf \frac{\tilde{\mu}_k}{\bar{\mu}^{n(k)}} \\ &\geq (\bar{\mu}\bar{\pi})^{t-1} \lim \frac{\tilde{\mu}_k/\bar{\mu}^{n(k)}}{\tilde{\pi}_k/\bar{\pi}^{n(k)}} \liminf \frac{\tilde{\pi}_k}{\bar{\pi}^{n(k)}} > 0, \text{ for all } t \in \mathbb{N}, \end{aligned}$$

where the last passage follows from noticing that $\lim \frac{\tilde{\mu}_k/\bar{\mu}^{n(k)}}{\tilde{\pi}_k/\bar{\pi}^{n(k)}} > 0$ for all equipoised sequences.

Part 2. Challenger is more committed: $\frac{\tilde{\mu}_k}{\bar{\mu}^{n(k)}} \in (\bar{\mu}, 1)$. Thus $\liminf \frac{\tilde{\mu}_k}{\bar{\mu}^{n(k)}} \geq \bar{\mu}$. By Proposition 1, for any $k \in \mathbb{N}$, the equilibrium probability of attack at time $t \in \mathbb{N}$ is given by

$$\Pr(\text{attack at } t) = \bar{\mu}^{t-1}\bar{\pi}^{t-2}\frac{\tilde{\pi}_k}{\bar{\pi}^{n(k)}}, \text{ for all } t \in [2, n(k)+1].$$

$$\begin{aligned} \therefore \liminf \Pr(\text{attack at } t) &= \bar{\mu}^{t-1}\bar{\pi}^{t-2} \liminf \frac{\tilde{\pi}_k}{\bar{\pi}^{n(k)}} \\ &\geq \bar{\mu}^{t-1}\bar{\pi}^{t-2} \lim \frac{\tilde{\pi}_k/\bar{\pi}^{n(k)}}{\tilde{\mu}_k/\bar{\mu}^{n(k)}} \liminf \frac{\tilde{\mu}_k}{\bar{\mu}^{n(k)}} > 0, \text{ for all } t \geq 2, \end{aligned}$$

where the last step follows from noticing that $\lim \frac{\tilde{\pi}_k/\bar{\pi}^{n(k)}}{\tilde{\mu}_k/\bar{\mu}^{n(k)}} > 0$ for all equipoised sequences.

□

A.4 Proofs of Section 4

Proof of (11)

Proof. From the definition of order of conflict, when Challenger is relatively more committed, then

$$n \approx \frac{\ln \mu_0}{\ln \bar{\mu}}.$$

From Proposition 1, upon observing an attack, Defender concedes with total probability

$$\frac{\pi_0}{\bar{\pi}^n} \approx \pi_0 \mu_0^{-\frac{\ln \bar{\pi}}{\ln \bar{\mu}}}.$$

Using Lemma 1,

$$\frac{\ln \bar{\pi}}{\ln \bar{\mu}} = \frac{\ln\left[1 - \left(1 - \delta^C\right) A/c\right]}{\ln\left[\delta^D\left(1 + \left(1 - \delta^D\right) L/d\right)\right]}.$$

Replacing δ with $e^{-r\Delta}$, differentiating with respect to Δ, and taking the limit as Δ goes to 0 we get

$$\lim_{\Delta\downarrow 0}\frac{\ln\bar{\pi}}{\ln\bar{\mu}}=\frac{\tilde{A}r^C/\tilde{c}}{r^D\left(\tilde{L}-\tilde{d}\right)/\tilde{d}}=\frac{\sigma^D}{\sigma^C},$$

and therefore

$$P^D=1-\frac{\pi_0}{\bar{\pi}^n}=1-\pi_0\mu_0^{-\frac{\sigma^D}{\sigma^C}}$$

which concludes the proof. It is worth noticing that σ^D and σ^C are the constant hazard rate of concession of Defender and Challenger in the limit case when $\Delta\downarrow 0$. □

Proof of Proposition 4

Proof. We first prove the proposition for the continuous-time case ($\Delta\to 0$). We then show that the solution to the model with discrete but frequent enough opportunistic to concede is the same as the solution of the continuous-time case.

Let $\alpha(x)$ be the probability with which an offer $x\in[0,1]$ by Defender is accepted by normal Challenger. Then, the total probability of acceptance is $\beta(x)=(1-\hat{\mu})\,\alpha(x)$. Let $\rho^i:=e^{-r^iT}$ denote the total discounting from the offer to the start of conflict (from $-T$ to 0). Challenger compares the offer with the payoff of triggering conflict, which itself depends on the offer. If Challenger accepts the offer x with probability $\beta(x)$, she gets $x\tilde{c}/r_C$ from accepting, whereas rejecting x raises the probability of the tough type of Challenger to $\mu_0=\hat{\mu}/(1-\beta(x))$. Therefore the value of conflict is given by

$$\rho^C P^D(\mu_0,\pi_0)\frac{\tilde{c}}{r^C}=\rho^C P^D\left(\frac{\hat{\mu}}{1-\beta},\pi_0\right)\frac{\tilde{c}}{r^C},\tag{19}$$

where the probability $P^D(\mu_0,\pi_0)$ with which Defender concedes as soon as conflict begins in the limit as $\Delta\downarrow 0$ is given by (11). Therefore, Defender's expected discounted share of the resource, as measured from date $-T$, as a function of his offer x and the total acceptance probability β, is

$$V(x):=\beta(1-x)\frac{\tilde{d}}{r^D}+(1-\beta)\left\{\hat{\pi}\frac{\tilde{d}}{r^D}+(1-\hat{\pi})\left(1-\rho^D\right)\frac{\tilde{d}}{r^D}\right\}.$$

Defender's problem is to choose the offer share x to get the maximum discounted utility

$V(x)\frac{\tilde{d}}{r^D}$ subject to Challenger wanting to accept the offer:

$$\max_{x\in[0,1]} \quad V(x)$$
$$\text{s.t.} \quad x \leq \quad \rho^C\left\{1-\hat{\pi}\left(\frac{\hat{\mu}}{1-\beta}\right)^{-(\sigma^D/\sigma^C)}\right\} \tag{20}$$

where $\beta = 0$ if the inequality is strict (because an offer that is worse than the outside option of conflict would not be accepted by Challenger with strictly positive probability).

Assuming that $\beta > 0$, so that the constraint holds as an equality, we can substitute for x in terms of β and take the total derivative of V with respect to β:

$$\frac{dV}{d\beta} = -x(\beta) - \beta\frac{dx}{d\beta} + \rho^D(1-\hat{\pi}),$$

where

$$\frac{dx}{d\beta} = \rho^C\hat{\pi}\left(\frac{\hat{\mu}}{1-\beta}\right)^{-(\sigma^D/\sigma^C)}\frac{1}{1-\beta}\cdot\frac{\sigma^D}{\sigma^C}.$$

Substituting,

$$\begin{aligned}\frac{dV}{d\beta} &= -\rho^C + \rho^C\hat{\pi}\left(\frac{\hat{\mu}}{1-\beta}\right)^{-(\sigma^D/\sigma^C)}\left(1-\frac{\beta}{1-\beta}\frac{\sigma^D}{\sigma^C}\right) + \rho^D(1-\hat{\pi}),\\ &= -\rho^C + \rho^C\hat{\pi}\left(\frac{\hat{\mu}}{1-\beta}\right)^{-(\sigma^D/\sigma^C)}\left(1-\frac{\beta}{1-\beta}\frac{\sigma^D}{\sigma^C}\right) + \rho^C(1-\hat{\pi}) + \left(\rho^D-\rho^C\right)(1-\hat{\pi}).\end{aligned}$$

Let V^e denote the function V in the special case when both discount factors are ρ^C . We then have:

$$\frac{dV^e(\beta)}{d\beta} = \hat{\pi}\rho^C\left\{\left(\frac{\hat{\mu}}{1-\beta}\right)^{-\sigma^D/\sigma^C}\left(1-\frac{\beta}{1-\beta}\frac{\sigma^D}{\sigma^C}\right)-1\right\}.$$

We know that

$$\left.\frac{dV^e}{d\beta}\right|_{\beta=0} = \hat{\pi}\rho^C\left\{\hat{\mu}^{-\sigma^D/\sigma^C}-1\right\} > 0, \tag{21}$$

$$\left.\frac{dV^e}{d\beta}\right|_{\beta=1-\hat{\mu}} = -\hat{\pi}\rho^C\left(\frac{1-\hat{\mu}}{\hat{\mu}}\frac{\sigma^D}{\sigma^C}\right) < 0. \tag{22}$$

Therefore:

$$\frac{dV}{d\beta} = \frac{dV^e}{d\beta} + \left(\rho^D-\rho^C\right)(1-\hat{\pi}). \tag{23}$$

Notice that $dV/d\beta$ is decreasing in β. Therefore, normal Challenger accepts the equilibrium offer is with probability less than 1 ($\beta < 1 - \hat{\mu}$) if

$$\left.\frac{dV}{d\beta}\right|_{\beta=1-\hat{\mu}} < 0.$$

Substituting (21) into (23) and rearranging terms we have

$$\left.\frac{dV}{d\beta}\right|_{\beta=1-\hat{\mu}} < 0 \iff \left(e^{-\left(r^D - r^C\right)T} - 1\right)\frac{r^D}{r^C} < \frac{\hat{\pi}}{1-\hat{\pi}}\frac{1-\hat{\mu}}{\hat{\mu}}\frac{\sigma^D}{\sigma^C}$$

where

$$\frac{\sigma^D}{\sigma^C} = \frac{\tilde{A}/\tilde{c}}{\left(\tilde{L} - \tilde{d}\right)/\tilde{d}}.$$

This proves the first statement of the proposition. The second statement follows immediately from noticing that

$$\lim_{T\to 0}\left(e^{-\left(r^D - r^C\right)T} - 1\right)\frac{r^D}{r^C} = 0 < \frac{\hat{\pi}}{1-\hat{\pi}}\frac{1-\hat{\mu}}{\hat{\mu}}\frac{\sigma^D}{\sigma^C}.$$

We now show that that the solution to the model with discrete but frequent enough opportunities to concede is the same as the solution to the continuous-time model. Let $V(x, \beta, \delta)$ denote the total discounted (to time $-T$) utility of the uninformed Defender when his offer of x is then accepted with probability β, and the length of time that elapses between successive opportunities to concede should conflict start:

$$V(x, \beta, \Delta) := \beta(1 - x) + (1 - \beta)\left(1 - \rho^D + \hat{\pi}\rho^D\right) - (1 - \beta)\left(1 - \delta^D\right). \tag{24}$$

Defender's problem is

$$(\mathcal{P}) \quad \max\ V(x, \beta, \Delta)\ \text{s.t.}\ (x, \beta) \in F(\Delta), \tag{25}$$

;where $F : [0, 1] \to [0, 1 - \mu_0] \times [0, 1]$ is the set of feasible pairs of (β, x). A value of Δ determines $\delta^i = e^{-r^i\Delta}$, which in turn determines the threshold $\bar{\mu}$ and $\bar{\pi}$ according to (4) and (5), which in turn determines the order of conflict $n(\Delta)$ as

$$\frac{\hat{\mu}/(1-\beta)}{\bar{\mu}^{n(\Delta)-1}} < \bar{\mu} \leq \frac{\hat{\mu}/(1-\beta)}{\bar{\mu}^{n(\Delta)}}. \tag{26}$$

This value of $n(\Delta)$ in turn determines the probability P^D with which Defender concedes at the start is given by (see (11))

$$1-\frac{\hat{\pi}}{\bar{\pi}^{n(\Delta)+1}}\leq P^{D}\left(\Delta,\beta\right)\leq 1-\frac{\hat{\pi}}{\bar{\pi}^{n(\Delta)}} \tag{27}$$

if the inequality defining n holds as an equality, and as

$$P^{D}\left(\Delta,\beta\right)=1-\frac{\hat{\pi}}{\bar{\pi}^{n(\Delta)}}; \tag{28}$$

each value of P^D determines a unique value of the offer $x\left(\Delta,\beta\right)$ that leaves Challenger indifferent between accepting and rejecting the offer:

$$P^{D}\left(\Delta,\beta\right)\cdot\frac{\tilde{c}}{r^{C}}-\tilde{A}\frac{1-\left(1-e^{-r^{C}\Delta}\right)}{r^{C}}=e^{r^{c}T}\cdot x\left(\Delta,\beta\right)\frac{\tilde{c}}{r^{C}}. \tag{29}$$

We assert that F is a compact-valued correspondence continuous at $\Delta=0$. The image of Δ under F can be written as a correspondence $\phi:\beta\twoheadrightarrow\phi(\beta)$ where $\phi\left(\beta\right)$ is a singleton given by (28) if (Δ,β) satisfies A2, and is otherwise a closed interval given by (27). Fix any β. Suppose $\left(\Delta^{k}\right)_{k\geq 1}$ is any sequence decreasing to 0; let $\left(\beta,x^{k}\right)_{k}$ be a sequence of points in $F\left(\Delta^{k}\right)$. Then from the construction of P^D we see that

$$P^{D}\left(\Delta,\beta\right)\rightarrow P^{D}\left(0,\beta\right),$$

which using (29) implies that $x^k\rightarrow x$. In other words, $(\beta,x)\in F\left(0\right)$. Since ϕ is uniformly bounded this means that if $\left(\beta^{k},x^{k}\right)_{k}$ is a sequence of points in $F\left(\Delta^{k}\right)$ converging to (β,x) then we have $(\beta,x)\in F\left(0\right)$. Therefore, F is upper hemicontinuous.

Suppose $(\beta,x)\in F\left(0\right)$ and $\left(\Delta^{k}\right)_{k\geq 1}$ is any sequence decreasing to 0. Then pick any sequence $\left(\beta^{k},x^{k}\right)_{k}$ such that $\beta^{k}=\beta$ for all k and $\left(\beta,x^{k}\right)\in F\left(\Delta^{k}\right)$. By construction the bounds in (27) and the expression in (28) reduce to $P^{D}\left(0,\beta\right)$ as $\Delta^{k}\rightarrow 0$. Thus it must be the case that $P^{D}\left(\Delta^{k},\beta\right)\rightarrow P^{D}\left(0,\beta\right)$, and therefore, by the continuity of (29), we have $x^{k}\rightarrow x$; hence F is lower hemicontinuous.

The function V is jointly continuous in the variables (x,β) and the parameter Δ. Since F is both upper hemicontinuous and lower hemicontinuout, which implies that F is continuous at $\Delta=0$, the maximum theorem immediately implies that the optimal solutions are upper hemicontinuous at Δ. In other words, all optimal solutions at a sufficiently small Δ are very close to the set of solutions at 0. □

Proof of Proposition 6

Proof. Let $x_{\mathtt{t}}$ and $x_{\mathtt{n}}$ denote the equilibrium offers of the two types. We know that $x_{\mathtt{t}} = 0$ in any equilibrium. If normal Defender offers any amount other than 0 he reveals his type and therefore Challenger will accept only such offers whose discounted value is no less than waiting until time 0 to attack and then extracting an immediate concession. The smallest share that would be accepted by Challenger is the discounted value of the entire pie:

$$x_{\mathtt{n}}^{\min}\frac{\tilde{c}}{r} = \rho^C\frac{\tilde{c}}{r} \Rightarrow x_{\mathtt{n}} \geq x_{\mathtt{n}}^{\min} = \rho^C.$$

The normal Defender's value from conflict, as a function of the probability q with which normal Challenger attacks after receiving a zero offer, is

$$U^D(q) := \left(1-\rho^D\right)\cdot 1 + \rho^D\left\{\hat{\mu}\cdot 0 + (1-\hat{\mu})\cdot(1-q)\right\},$$

where the probability q naturally depends on the equilibrium under consideration.

In the separating equilibrium, $x_{\mathtt{n}} \geq x_{\mathtt{n}}^{\min} = \rho$, $x_{\mathtt{t}} = 0$, and therefore $q = 0$ by Bayes' rule (only tough Defender offers zero). So the normal Defender's incentive compatibility condition is

$$1 - x_{\mathtt{n}} \geq U^D(0) \Longrightarrow x_{\mathtt{n}} \leq \rho_D\hat{\mu}.$$

In the pooling equilibrium, $q = 1$ since the prior is unchanged. Normal Defender does not deviate and offer $x_{\mathtt{n}}^{\min}$ if

$$1 - x_{\mathtt{n}}^{\min} \leq U^D(1) \Longrightarrow \rho^C \geq \rho^D.$$

In the semi-pooling equilibrium, Defender is indifferent if

$$1 = \left(1-\rho^D\right)\cdot 1 + \rho^D\left\{\hat{\mu}\cdot 0 + (1-\hat{\mu})\cdot(1-q)\right\}$$

and hence

$$\hat{\mu} + (1-\hat{\mu})\,q = \frac{\rho^C}{\rho^D}.$$

A value of q satisfying the above exists if and only if ρ^C/ρ^D lies in $[\hat{\mu}, 1]$. If $q > 0$ then it must be the case that the normal type of Defender mimics the tough type (by offering zero) with a high enough probability that the posterior after a zero offer still makes it profitable for Challenger to attack. $\square$

A.5 Proofs of Section 5

Proof of Proposition 7

Proof. Let the posterior probability of the tough type at the end of round k be $\hat{\mu}_k$. A strategy of Defender comprises finitely many functions x_k for $k = 0, 1, \ldots, K$, mapping from $[0,1]$ to $[0,1]$ such that the k^{th} offer is $x_k(\hat{\mu}_{k-1})$. Let $(\beta_k \mid k = 0, 1, \ldots, K)$ be the corresponding 'acceptance probability' mappings from $[0,1]$ to $[0,1]$; the offer made in round $k \in \{0, 1, \ldots, K\}$ is accepted with probability $\beta_k(\hat{\mu}_{k-1})$. The strategy of Challenger decides which offer to accept and with what probability, one for each round; we omit the notation for this.

We prove this by induction on the number of rounds. Let $K = 2$ henceforth. The optimal offer at each state and each history is (i) Markovian, i.e. it depends only on beliefs about the type of Challenger; and (ii) deterministic. If B is the function such that $\beta^* = B(\hat{\mu})$ is the unique solution to the equation (14), then it is clear that B is increasing in $\hat{\mu}$ (strictly increasing unless it has hit 1), and so is the posterior probability $\hat{\mu}/(1 - B(\hat{\mu}))$. Take any candidate equilibrium of the 2-offer game with the equilibrium acceptance functions (β_1, β_2). Clearly, sequential rationality requires that $\beta_2(\hat{\mu}_1) = B(\hat{\mu}_1)$, where $\hat{\mu}_1$ is given by $\hat{\mu}_1 = \hat{\mu}/\left(1 - \bar{\beta}_0\right) > \hat{\mu}$. The total probability (summed over rounds and types) that conflict will not start is then given by $\bar{\beta} = \beta_1 + (1 - \beta_1)\beta_2$. The posterior at the end of round 2 is

$$\mu_0 = \frac{\hat{\mu}_1}{1 - \beta_2} = \frac{\hat{\mu}}{(1 - \beta_1)(1 - \beta_2)} = \frac{\hat{\mu}}{1 - \bar{\beta}}.$$

By the above and equation (**??**), the sequentially rational offer in round 2 is

$$\bar{x} = e^{-r\Delta_{1,2}} \left\{ 1 - \hat{\pi} \left(\frac{\hat{\mu}}{1 - \bar{\beta}} \right)^{-\sigma^D/\sigma^C} \right\} \tag{30}$$

where $\Delta_{1,2} \geq 0$ is the time interval between the first and the second offer. Incentive compatibility of normal Challenger then requires that acceptance in any round give the same utility to Challenger. This, together with the fact that $\left(\bar{x}, \bar{\beta}\right)$ satisfies (30), means that the utility of Defender is the utility of a game in which he takes out a mass $\bar{\beta}$ of Challengers (at a cost of $\bar{x}$ per unit mass); this is denoted by $V_0^D\left(\bar{x}, \bar{\beta}; \hat{\mu}, \hat{\pi}\right)$. Let $(x^*(B(\hat{\mu})), B(\hat{\mu}))$ be the optimal pair when we have a single round with the prior $\hat{\mu}$. Since $\bar{x} > x^*$ and

$\bar{\beta} > B(\hat{\mu})$, it follows that

$$V_0^D\left(\bar{x}, \bar{\beta}; \hat{\mu}, \hat{\pi}\right) < V_0^D\left(x^*\left(B\left(\hat{\mu}\right)\right), B\left(\hat{\mu}\right); \hat{\mu}, \hat{\pi}\right).$$

If Defender deviates in round 1 from β_1 and chooses $\beta_1' = 0$ instead, it would have been sequentially rational to make the optimal offer x^* in round 2 and have it accepted with probability β^*. Therefore, if two offers are accepted with positive probability in any equilibrium, it is better to deviate and change the earlier offer to 0, so that the first offer is rejected. Hence there is no such equilibrium: the only possibility is $\beta_1^* = \ldots = \beta_{K-1}^* = 0$. □