Figure 2: Sample equilibrium path where player 2 wins

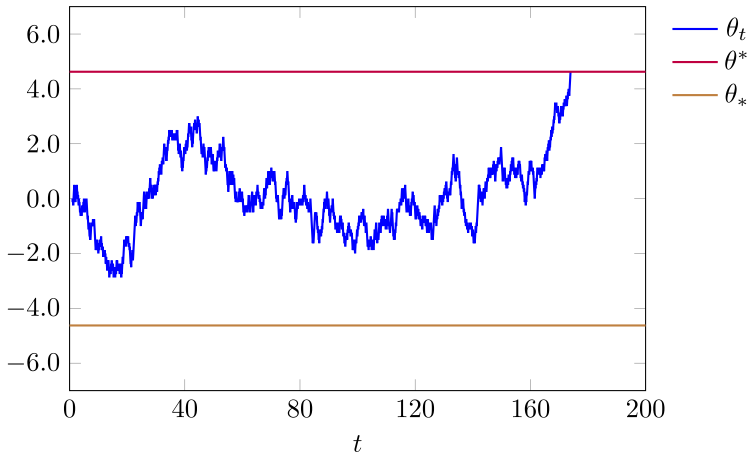

over time, which structurally favors player 2, whence both thresholds are much lower. Figure 2 shows an example of a typical equilibrium path, under the same parameter assumptions as Figure 1a and starting at $\theta_0 = 0$: the state of the world is initially between the two thresholds, so both players continue the war until θ reaches one of them, in this case player 1's.

5 Equilibrium in the Limit

When applying the model, the above analysis is relevant if we consider the movement of θ to be a significant feature of the application. However, the model also admits being used simply as an equilibrium selection tool when the "true" model the researcher is interested in is the basic war of attrition without perturbations. In this interpretation, we must take the limit of the solution as the movement of θ becomes arbitrarily slow. Formally, we apply the same transformations as in the previous Section to the stochastic process: the density of $\theta' - \theta | \theta$ is now $\frac{1}{\nu} F_\theta \left(\frac{x}{\nu} + \mu(\theta)(1-\nu) \right)$, with mean $\mu(\theta)\nu^2$ and variance $\sigma^2\nu^2$; but the costs are left unchanged, and we take the limit as $\nu \to 0$. In this case the fundamentals of the game are no longer the same, as the movement of θ has been slowed relative to the size of the flow costs.

Naturally, the limit of the game as $\nu \to 0$ is simply the unperturbed war of attrition, but taking the limit of the perturbed equilibria yields a uniquely selected equilibrium of the basic game:

Proposition 4 (Equilibrium Selection with Slow-Moving Processes)**.** *Suppose* $H_1 = H_2 =$

H.[8] *Let θ^l be such that $c_1(\theta^l) = c_2(\theta^l)$. Suppose $c_1'(\theta^l) = -c_2'(\theta^l)$. Let $\theta^*(\nu)$, $\theta_*(\nu)$ be the thresholds of the equilibrium in Proposition* **??** *as a function of ν. Then $\theta^*(\nu), \theta_*(\nu) \to \theta^l$ as $\nu \to 0$.*

Hence, if $\theta_0 < \theta^l$ the perturbed equilibria converge, as $\nu \to 0$, to an equilibrium of the unperturbed game where 1 wins instantly. If $\theta_0 > \theta^l$, then 2 wins instantly in the limit.

If $\theta_0 = \theta^l$ the perturbed equilibria converge to an equilibrium of the basic war of attrition augmented with tokens. *In the augmented game, players observe a payoff-irrelevant variable ϕ_t that follows a random walk, starting at 0 and obeying $\phi_{t+1} = \phi_t + 1$, $\phi_{t+1} = \phi_t - 1$ with probabilities 0.5, 0.5. The limiting equilibrium is described by a unique threshold K such that 1 surrenders at any history where $\phi_t \geq K$, and 2 surrenders whenever $\phi_t \leq -K$.*

In this equilibrium, both players have probability 0.5 of winning ex ante, and there is delay, but both players are strictly willing to fight, and i's expected payoff ex ante is $\frac{H}{4} > 0$.

In other words, as the perturbation vanishes, the model offers a stark prediction for what equilibrium should be selected in the unperturbed game: if one player has a higher flow cost[9] than the other, then the equilibrium where she always surrenders immediately, and the other player is never expected to surrender, should be played. This is not surprising, since the perturbed game has a unique equilibrium, and any family of unique equilibria for the unperturbed game displaying the right comparative statics (i.e., higher-cost players should be more likely to lose) must make this selection. More interestingly, though, the symmetric equilibrium selected when costs are equal is *not* the totally mixed equilibrium of the unperturbed game, and in fact is not an equilibrium of the unperturbed game, as it requires a coordination mechanism. Indeed, the equilibrium is *as if* each player were initially in possession of K tokens, and a stage game were played every period; both players would have equal probability of winning in each stage game, and the loser would pay one token to the winner. (In terms of the formal statement, $\phi_t + K$ is the amount of tokens held by player 2 at time t.) Then the selected equilibrium prescribes that a player should surrender whenever she runs out of tokens.

6 Discussion

In this Section we discuss the merits of the model's results and compare them to existing models of wars of attrition, namely the basic unperturbed model and an alternative based on reputational concerns. Although these alternatives are not novel, for completeness we briefly define the models and state their basic results; the details are in Appendix A. We

[8] This serves to simplify notation.

[9] In general, since the benefits from winning may differ, what matters is the cost-to-benefit ratio.

focus on the continuous time versions of both models as the discrete time versions introduce some technical complications which are not relevant to our purposes.[10]

6.1 Basic War of Attrition

The basic, unperturbed war of attrition corresponds to a special case of our model where $\mu = \sigma = 0$, i.e., θ is constant. Alternatively, we can take $c_1(\theta) \equiv c_1$ and $c_2(\theta) \equiv c_2$ to be flat. The following Proposition summarizes its set of SPE:

Proposition 5 (Equilibria of the Basic War of Attrition)**.** *The basic war of attrition has three SPE:*

- *An equilibrium where* 1 *surrenders at every history and* 2 *never surrenders.*
- *The opposite equilibrium where* 1 *wins immediately.*
- *A totally mixed equilibrium where both players mix at every history, and* i *chooses to surrender at a rate* $p_i = \frac{c_j}{H}$ *at every* t*. Players are indifferent about continuing. Players' expected payoffs ex ante are* 0*.* i*'s probability of winning is* $\frac{c_i}{c_i+c_j}$*.*

As noted in the Introduction, this model has two main shortcomings. First, it provides no way to select an equilibrium from the options offered, which saps it of explanatory power as the set of equilibria includes extremes where either player wins for sure. Second, the mixed equilibrium seems reasonable in the symmetric case, but it clearly has dysfunctional comparative statics away from it: indeed, a player's probability of winning increases with her *own* cost, and in the limit where one player has much higher costs than the other, that player wins almost surely. These results derive from the mechanics driving the equilibrium: since both players must be indifferent along the equilibrium path, if i has high costs, then j must surrender at a high rate to keep i willing to continue with positive probability. This has little to do with the natural intuition behind the game, namely, that if i knows j has lower costs, she might conclude that she probably cannot convince j to concede, leading to i's surrender.

The perturbed model overcomes these issues: the movement of θ results in a unique equilibrium being selected. Moreover, changes in the players' costs are parameterized simply by movements in θ, which change the payoffs and probabilities of winning in the natural way, i.e., the player with higher costs is more likely to surrender sooner and vice versa.

[10] Namely, the discrete time basic war of attrition has a partially mixed equilibrium in addition to the totally mixed one, where players take turns mixing vs. continuing for sure; this equilibrium converges to the totally mixed equilibrium in the continuous time limit. Similarly, the equilibrium of the war of attrition with reputation involves alternation in mixing.

In addition, the equilibrium selected by the perturbed game in the symmetric case (or, away from the slow-movement limit, in other cases where the war continues for some time) produces strictly better welfare than the totally mixed equilibrium of the basic game. As per Proposition 4, the sum of the players' expected payoffs is $\frac{H}{2}$, whence the expected delay is $\frac{H}{4c}$, versus a total payoff of 0 and expected delay $\frac{H}{2c}$ in the totally mixed equilibrium. It should be noted, though, that the equilibrium selected by the perturbed model requires a coordination device, and if such devices are allowed then even more efficient equilibria are possible. Indeed, the most efficient equilibrium possible would involve the players flipping a coin, with the loser of the coin flip conceding immediately. This would generate a total payoff of H with no delay. It is an interesting fact in its own right, though, that this equilibrium *cannot* be selected by the perturbed model.

6.2 War of Attrition with Reputation

The war of attrition with reputation takes the basic war of attrition outlined above, but adds for each player i a probability ϵ_i of being a *commitment type* that never surrenders. Types are private information. The following Proposition characterizes the SPE of this game:

Proposition 6 (Equilibrium of the War of Attrition with Reputation)**.** *The war of attrition with reputation has a unique SPE, characterized as follows. Let t_i^* be the* vanishing time *of player i, characterized by $e^{\frac{-c_j}{H}t_i^*} = \epsilon_i$. If $t_i^* = t_j^*$, then we say the players are* balanced*. The probability that i (rational or not) will continue up to time t is $Q_{it} = e^{\frac{-c_j}{H}t}$ for $t \leq t_i^*$. The mass of rational i's surviving up to time t is $P_{it} = e^{\frac{-c_j}{H}t} - \epsilon$. The surrender rate for rational i's is $p_{it} = \frac{c_j}{H}\frac{1}{1-\epsilon e^{\frac{c_j}{H}t}}$; in particular, rational players are always mixing. At time $t_i^* = t_j^*$, only commitment types are left and there is no surrendering thereafter.*

If $t_i^* < t_j^*$, we say that i is stronger than *j. In this case, a mass of j's rational types of size $P_{j0} = 1 - \frac{\epsilon_j}{\epsilon_i^{\frac{c_i}{c_j}}}$ surrender immediately at $t = 0$; after that, the players are balanced and the equilibrium continues as above. In this case, the probability of winning for i's and j's rational types are*

$$W_i = P_{j0} + \frac{1-P_{j0}}{1-\epsilon_i}\left[\frac{c_i}{c_i+c_j}\left(1 - e^{-\frac{c_i+c_j}{H}t_i^*}\right) - \epsilon_i\left(1 - e^{-\frac{c_i}{H}t_i^*}\right)\right] \approx P_{j0} + (1-P_{j0})\frac{c_i}{c_i+c_j},$$
$$W_j = \frac{1-P_{j0}}{1-\epsilon_j}\frac{c_j}{c_i+c_j}\left(1 - e^{-\frac{c_i+c_j}{H}t_i^*}\right) - \frac{\epsilon_j}{1-\epsilon_j}\left(1 - e^{-\frac{c_j}{H}t_i^*}\right) \approx (1-P_{j0})\frac{c_j}{c_i+c_j}.$$

i's expected payoff is $P_{j0}H$; j's expected payoff is 0*.*

The mechanics behind the equilibrium are as follows. After $t = 0$, rational types must be willing to mix, as contradictions arise in all other cases. As in the unperturbed model, for i to be willing to mix, j must be surrendering continuously at a rate $\frac{c_i}{H}$ (this is j's overall surrender rate regardless of her rationality, so j's rational types must be surrendering more often than this value to maintain the correct average), and vice versa. At this rate, the rational types of player i would all surrender by time t_i^*. If these vanishing times were unequal, and e.g. $t_i^* < t_j^*$, then j's rational types would surrender at time t_i^*, which would incentivize i's rational types in some interval $[t_i^* - \epsilon, t_i^*]$ *not* to surrender, and so on. To prevent this, a mass of j's rational types surrender at the beginning so that in the continuation the rational types of both players will extinguish themselves at the same time.

As in the unperturbed war of attrition, the higher-cost player fares better (in the sense of surrendering less often) *after* $t = 0$. However, if j has a higher flow cost, that also results in a higher t_j^*, and therefore a higher probability of surrendering immediately. Overall the latter effect wins, so that having a higher flow cost reduces a player's equilibrium payoff and probability of winning. However, the vanishing times are also a function of the ϵ_i, so guaranteeing that a certain relationship between flow costs translates into a certain player likely winning the war requires some assumption on these parameters. Specifically, if $c_i < c_j$, then as ϵ_i, $\epsilon_j \to 0$ we have that $P_{j0} \to 1$, but only if $\frac{\epsilon_j^{c_j}}{\epsilon_i^{c_i}} \to 0$. In other words, ϵ_i cannot be going to 0 much faster than ϵ_j; otherwise, this effect would overrule the relationship between the flow costs and j would be favored to win the war. When applying the model, we might be willing to assume that $\epsilon_i \approx \epsilon_j$ so that there is no issue, but if the ϵ's are expected to be very small this assumption will be hard to verify or disprove based on any sort of data the researcher has access to.

In this aspect, the model proposed in this paper appears more robust: if we are interested in using the perturbation as an equilibrium selection tool, the results as $\nu \to 0$ do not depend on the shape of the functions $c_i(\theta)$ away from the θ_0 we are focused on. They only depend on $(c_1(\theta_0), c_2(\theta_0))$ as shown by Proposition 4, so long as the functions are continuous.

Two other interesting differences should be noted, which might be used to distinguish between the models empirically. First, in the symmetric case, the reputational model again predicts the same totally mixed equilibrium as the basic war of attrition, unlike the moving-state model. Second, the density functions predicted for the length of the war $h(t)$ differ substantially between the two models, as shown in Figure 3. To begin with, the totally mixed equilibrium results in a higher expected delay until the war ends. But, more importantly, in the case of a relatively balanced war, the moving-state model predicts that there is no chance of an immediate surrender, and the rate of surrender is hump-shaped over time: on average it will take some time until θ moves enough to one of the extremes so as to make one player surrender. On the other hand, the reputational model predicts a substantial chance of immediate surrender, followed by exponential decay.

Figure 3: Density of length of the war

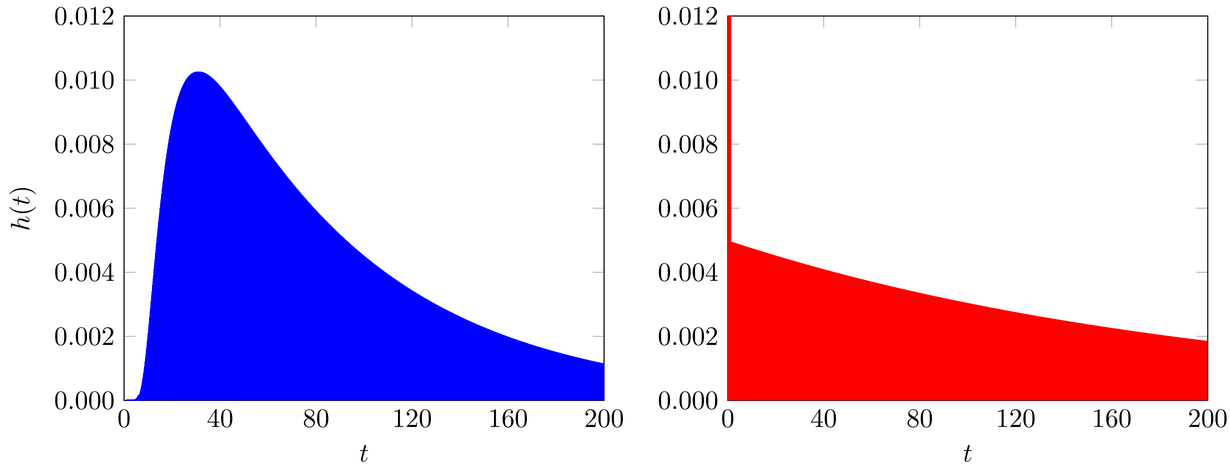

(a) Moving θ

(b) Reputation

7 Partial Concessions

We now extend the model to study the following variant of the war of attrition: as before, players are able to continue the war until one surrenders, but while the war is ongoing they can affect the payoffs of the eventual outcome by making partial concessions. For instance:

- An army sieges a city. The city has strong walls and cannot be taken by assault, so a siege will lead to two outcomes: either the army surrenders (meaning it leaves with nothing) or the city surrenders (meaning it opens its doors and the army plunders it). In this example, movements in θ reflect changes in the fortunes of each side: the city's supplies may dwindle or they may be able to smuggle in fresh supplies, either side may suffer a disease outbreak, and so on. Instead of surrendering outright, the city can instead pay tribute to the army: that is, it can gather some wealth, give it to the army and invite them to leave. If this tactic is successful, it is more efficient for both sides, since the city's gathering of its own wealth entails lower welfare losses (no buildings are burned, no civilians are killed, and so on). There is no commitment device, so the army can still stay to siege after receiving tribute. But the tactic may work because the remaining loot to be had is smaller, while the city is still eager to defend itself (since the deadweight loss from being plundered is large).

- A polluting firm is boycotted by an activist group. The boycott is costly for both players, as it lowers firm sales and consumer surplus. The war ends when either the firm capitulates to the demands or the activists abandon the boycott. However, the firm may have access to a range of policies it can implement to lower its own pollution.

A partial concession, in the form of a moderate level of self-regulation, may be enough to "deflate" the momentum of the boycott even if it is not what the activists demanded. Similar logic can apply to a government implementing emergency measures to appease protesters.

The common theme in these examples is that making a partial concession may benefit the party who does it: although it entails giving up some payoffs, this is often worth it if it will tilt the remaining war of attrition towards the side who partially concedes. We ask our model to answer two main questions: when are partial concessions useful? (For instance, are they only useful when they enable higher efficiency than the noncooperative outcomes, as in the siege example? Or can they be useful even when there is no efficiency gain involved?) And what is the optimal size and timing of a partial concession? We first study the case where only player 1 can make partial concessions, and then the case where both players can. In addition, we compare the results under two assumptions about timing: in one version, we take concessions to be one-shot, i.e., there is a fixed time when a concession must be made, or else the chance to do so is lost forever; in the other, multiple concessions can be made whenever the player wants, subject only to the restriction that concessions cannot be taken back.

One-Sided Concessions

The Limit Case

For simplicity and to fix ideas, we first discuss the limit case where $\nu \approx 0$. First, assume a one-shot concession at the beginning of the game: player 1 chooses a concession size x at $t = 0$. Thereafter, the game continues as usual, except that she gets $H - \alpha x$ from winning compared to losing, while 2 gets $H - x$ from winning. Here α represents the relative efficiency of the concession technology: if $\alpha < 1$, a concession reduces 1's incentive to win the war relatively less than 2's incentive, and vice versa.

Note that $\alpha < 1$ ($\alpha > 1$) does not necessarily imply that the concession generates (destroys) social welfare, as payoffs may have been normalized differently for each player. For instance, in the siege example, the benefits from winning before the concession would be higher for the defender than the attacker, i.e., $H_D > H_A$, since a losing defender suffers the cost of being plundered which is not collected by the attacker; and there would be some flow costs c_D, c_A. A transfer x from D to A would change D's payoffs to $H_D - x$ for winning or 0 from losing, while A's would change to H_A or x, respectively, generating no net change in welfare. However, if at the outset we had normalized the problem so that benefits were equal for both sides, then we might have normalized payoffs $H = \tilde{H}_D = \alpha H_D = \tilde{H}_A = H_A$;

$\tilde{c}_D = \alpha c_D$; $\tilde{c}_A = c_A$; and a transfer of x would change the benefits of winning to $H - \alpha x$ for D and $H - x$ for A, where $\alpha = \frac{H_A}{H_D}$.

Proposition 7 (One-Shot Concession with Slow-Moving θ). *The equilibrium of the game is as follows:*

- *If $c_1 < c_2$, player 1 makes no concessions and wins immediately.*
- *If $c_1 \geq c_2$ and $\alpha < 1$, player 1 makes a concession x^* defined by:*

$$\frac{H - \alpha x^*}{c_1} = \frac{H - x^*}{c_2}$$

  and wins immediately. This leads to a concession $x^* = H\frac{c_1 - c_2}{c_1 - \alpha c_2}$, and payoffs $U_1 = H - \alpha H \frac{c_1 - c_2}{c_1 - \alpha c_2}$, $U_2 = H\frac{c_1 - c_2}{c_1 - \alpha c_2}$.
- *If $\alpha \geq 1$, concessions never benefit player 1. Hence, if $c_1 = c_2$, no concession is made ($x^* = 0$) and the equilibrium described in Proposition 4 is played. If $c_1 \geq c_2$, any $x^* \in [0, H]$ is compatible with equilibrium but player 1 loses immediately in all cases.*

In other words, when θ moves very slowly, partial concessions can be useful–but only when they strengthen the conceding player's relative cost-benefit ratio, i.e., when $\alpha < 1$. In this case they are very powerful, and indeed guarantee that player 1 *always* wins the game in equilibrium, even when c_1 is substantially higher than c_2. Of course, this does not necessarily mean that player 1 walks away with a large payoff, as the concession needed to win may be large. In particular, the required concession $H\frac{c_1 - c_2}{c_1 - \alpha c_2}$ is small when $\alpha << 1$ (so that player 1 only needs to take a small dent in her own payoff to substantially reduce 2's incentive to win), or when c_1 and c_2 are close (so 1 is not very disadvantaged to begin with, so a small nudge is enough to make 2 surrender). Conversely, as $\alpha \to 1$, if $c_1 > c_2$ the required concession converges to H.

On the other hand, when $\alpha \geq 1$–that is, when a concession weakens 1's relative cost-benefit ratio–making a concession can never help player 1. This seems like a natural result, although we will see that matters are subtler away from the limit case.

Finally, note that in the limit case the timing of concessions actually does not matter: although we have stated the above result in a game where 1 can only make a concession at $t = 0$, the equilibrium is trivially the same if 1 is instead allowed to make multiple concessions at any time. Indeed, if $\alpha < 1$ or $c_1 < c_2$, so that 1 is able to win in equilibrium, she might as well make the winning concession immediately, so as to avoid any cost of waiting.

7.0.1 General Processes

We limit ourselves here to the case of one-shot concessions:

Proposition 8 (Equilibrium with a One-Shot Concession)**.** *The equilibrium of the game is as follows:*

- *If $\theta < \theta_*$, player 1 makes no concession and wins immediately.*
- *If $\theta \geq \theta_*$, let $(\theta_*(x), \theta^*(x))$ be the thresholds resulting in the continuation game after a concession of size x is made. Let x_1 be the smallest concession for which $\theta_*(x) \geq \theta$, i.e., player 1 wins immediately in the continuation (we take $x_1 = \infty$ if no x satisfies this requirement). Let $x_2 \in [0, x_1]$ be a minimal concession that maximizes 1's surrender threshold, i.e.,*

$$x_2 = \min \left(argmax_{x \in [0,x_1]} \theta^*(x) \right)$$

  Then x_2 is the optimal concession. Moreover, $\theta^*(x)$ is quasiconvex, so either $x_2 = 0$ or $x_2 = x_1$. In the first case, no concession is made; in the second, player 1 wins immediately.

In particular, keeping all other parameters fixed, there is $\alpha^* < 1$ such that for all $\alpha \geq \alpha^*$ it is optimal to choose $x^* = 0$, while for $\alpha < \alpha^*$ the optimal choice is $x^* = x_2$. In addition, player 1's equilibrium utility is decreasing in α.

The intuition behind this result is as follows. By Proposition 3, as long as θ is far from the surrender thresholds, a concession which lowers 1's potential payoff from winning–while also changing the surrender thresholds–actually has no impact on her expected utility except by changing said thresholds. This follows immediately from the fact that H does not feature in the expression for expected utility. Moreover, this expression is clearly increasing in θ^*, whence 1 effectively wants to maximize her own surrender threshold. Two surprising implications follow.

First, the optimal concession may not be one that results in an immediate win for player 1: indeed, she may choose to make a positive concession which still leaves the outcome of the war in doubt, and produces delay.

Second, whenever $\alpha > 1$ concessions always hurt player 1, but this is also true for some values of α that are below 1. The intuition behind this result is that there are two forces which affect the surrender thresholds when x changes, as anticipated in Proposition 2. On the one hand, a higher x reduces the stakes of winning for both players; this reduces their willingness to accept delay, so it tends to bring both thresholds closer to each other while having no clear impact on their average. On the other hand, if $\alpha < 1$ (> 1) the concession tends to increase (decrease) both thresholds and as it strengthens (weakens) player 1's relative strength. For a positive concession to be optimal, there must be a range of concession values for which the latter effect dominates the former, which requires α to be below 1 by a certain amount. Note that, in particular, this means that player 1 may decline to make a concession even when this might guarantee a win for a small cost: for example, if $\alpha = 1$ and θ is initially close to $\theta_*(0)$, a small concession might be enough to bring $\theta_*(x)$ up to θ, clinching the win; but, since this would also decrease θ^*, it follows that the cost of the concession is higher than the gains in terms of reduced delay and increased probability of winning.

Two-Sided Concessions

We now move on to a further extension where both players are able to make concessions. This can be thought of as a form of bargaining by means of unilateral offers. To fix ideas, consider the following examples:

- Two players, a buyer and a seller, engage in a war of attrition over the price at which B might buy a good worth H. As the war continues, B can make statements of the form: "I will buy for any price up to x", while S can say, "I will sell for any prize above y". These promises are binding, so B's acceptable price x_t cannot decrease over time and vice versa. Whenever a player surrenders, the transaction happens at the price set by the other player's most recent statement.

- Two armies contest a one-dimensional territory $[0, 1]$, where $v_i(t)$ is the value assigned to point t by player i, and $H_i = \int_0^1 v_i(t)dt$. Assume v_1 is decreasing and v_2 is increasing. While the war is ongoing, the armies station troops at outposts throughout the territory to contest its control and engage in intermittent fighting, generating flow costs. Along the way, each player i can make credible, irreversible commitments to pull her troops out of a measurable region $S \subseteq [0, 1]$. When a player surrenders, she keeps only the territory that the other player conceded, and the other player takes the rest.

As a reasonable compromise between generality and simplicity, we allow for each player i to have a *concession factor* α_i, so that if initial rewards from winning are H_i, H_j and i makes a concession x, rewards become $H_i - \alpha_i x$, $H_j - x$. In the bargaining example, we would have $\alpha_1 = \alpha_2 = 1$, since concessions take the form of transfers. In the war for territory, the α_i would vary depending on the exact territory conceded, but at the margin we would have $\alpha_1\alpha_2 < 1$: indeed, if i is trying to reduce j's incentives to fight relative to her own, the optimal strategy is to concede territory that j values most and i least, i.e., 1 would make concessions of the form $[t, 1]$ and 2 would make concessions of the form $[0, t']$, which result in the conceding player's incentive to win going down less than the recipient's.

We will focus on the limit as $\nu \to 0$. We normalize the costs and benefits so that $c = c_1 = c_2$ to streamline the notation. For clarity, we restrict the players to the following concession protocol: at the beginning of the game, 1 chooses a partial concession, then 2 observes it and makes a partial concession of her own; then the rest of the game proceeds as usual. (It can be shown that switching the order of concessions, adding more alternating concessions, making them simultaneous or allowing concessions after $t = 0$ have no impact on the equilibrium concessions and payoffs.)

Proposition 9 (Two-Sided Concessions with Slow-Moving Processes)**.** *The equilibrium of the game is as follows. First, if $\alpha_1 \geq 1$, $\alpha_2 \geq 1$ then concessions are not used and the equilibrium is as in Proposition 4. Similarly, if $\alpha_1 < 1$, $\alpha_2 \geq 1$ then only player 1 may use concessions and the equilibrium is as in Proposition 7. If $\alpha_1 < 1$, $\alpha_2 < 1$ then:*

- *If $H_1 < \alpha_1 H_2$, player 2 wins immediately without making a concession and payoffs are $(U_1, U_2) = (0, H_2)$.*[11]

- *If $H_2 < \alpha_2 H_1$, player 1 wins immediately without making a concession and payoffs are $(U_1, U_2) = (H_1, 0)$.*

- *Otherwise, both players make partial concessions $x_1 = \frac{H_2 - \alpha_2 H_1}{1 - \alpha_1 \alpha_2}$, $x_2 = \frac{H_1 - \alpha_1 H_2}{1 - \alpha_1 \alpha_2}$ which jointly exhaust the prize, and payoffs are*

$$U_1 = \frac{H_1 - \alpha_1 H_2}{1 - \alpha_1 \alpha_2}, \qquad U_2 = \frac{H_2 - \alpha_2 H_1}{1 - \alpha_1 \alpha_2}.$$

The intuition behind the result is as follows. As before, a player can try to use a concession to improve her relative strength and win the war. However, player 1 understands that making a concession just large enough so that $H_1 - \alpha_1 x_1 \geq H_2 - x_1$, as she would do in Proposition 7, is not enough here because player 2 will undercut her with her own concession. Hence, in equilibrium player 1 concedes enough so that player 2 cannot win except by conceding the entirety of the remaining prize. The resulting payoffs turn out to be the same as when player 2 concedes first and follows the same logic.

[11] Technically there is a collection of equilibria giving the same outcome, as player 1 can make concessions that have no impact on payoffs.

8 Conclusions

We have analyzed a model of wars of attrition with an evolving state of the world. As intended, the perturbation used in the model yields several attractive properties which are absent from the baseline model. In particular, the game has a unique equilibrium, yielding predictions about the likely winner which reflect a natural intuition about how this game might play out in practice. Moreover, the solution exhibits well-behaved comparative statics, meaning that if a player's benefit from winning increases or her flow cost decreases, she will be more likely to win. In the same vein, such a change will lead to the war ending sooner if the player being strengthened held an advantage to begin with, while the war will lengthen if she was an underdog at first. These properties are shared by models based on reputational perturbations, but unlike those, the present model is more robust to assumptions about how the game looks away from the perturbation, an attractive property for applied researchers who may not be able to observe enough detail about a real example to guarantee that they are modeling perturbations correctly.

In addition, the present model is highly tractable, which allows it to be used as a modeling tool in other settings or in applied problems. This is illustrated in Section 7, which shows how to account for the possibility of partial concessions during a war of attrition, and how access to this option might benefit the player making concessions. But many other applications are possible: for instance, the game can be extended to explore the use of costly commitment devices, i.e., "burning bridges"; to wars of attrition involving more than two players, as in legislative standoffs; or to cases where players have some control over the flow costs, such as price wars.

Finally, the model makes a novel prediction regarding the most natural equilibrium in balanced wars of attrition (where both players have similar chances of winning), which entails higher payoffs than a totally mixed equilibrium, as well as a different distribution for the expected length of the war. This raises the question whether this equilibrium is indeed played in real-life scenarios, or under what circumstances (e.g., does there have to be a changing, publicly-observed variable which is focal for the players?), a question which might be answered by laboratory or field experiments.

A Appendix

[[[add measurability to strategies]]]

Definition 1. A (non-terminal) history at time t, h_t, is a sequence of states of the world $(\theta)^t = (\theta_0, \dots, \theta_t)$, and the sequence of actions $(a_{is})_{i,s}$ given by $a_{is} = 1$ for $i = 1, 2$ and $s = 0, \dots, t$.

We will say histories to refer to non-terminal histories for brevity, and write $h_t = (\theta)^t$ as shorthand for $h_t = ((\theta)^t, (1, \dots, 1), (1, \dots, 1))$.

Definition 2. A strategy $\sigma_i(h)$ for player i is *Markov* if $\sigma_i(h) = \sigma_i(h')$ for all $h = (\theta)^t$, $h' = (\theta)^{t'}$ such that $\theta_t(h) = \theta_{t'}(h')$.

A strategy σ_1 for player 1 (2) is a *threshold strategy* with threshold θ^* if $\sigma_1(h_t) = 1$ whenever $\theta_t(h_t) < \theta^*$ ($>$) and $\sigma_1(h_t) = 0$ whenever $\theta_t(h_t) > \theta^*$ ($<$).

Lemma 1. *In any SPE, player* 1 *never surrenders at time* t *if* $\theta_t < -M_1$*, and player* 2 *never surrenders at time* t *if* $\theta_t > M_2$*.*

Proof. For player 1, surrendering when $\theta_t < -M_1$ yields a continuation payoff of 0, while continuing until the first time $s > t$ when $\theta_t \geq -M_1$ yields a strictly positive payoff. The proof for player 2 is identical. □

Definition 3. For any history h and strategies σ_i, σ_j, let $U_i(\sigma_i, \sigma_j | h)$ be i's continuation utility from strategy profile (σ_i, σ_j) at history h.

Definition 4. For any history h_0, any strategy σ_i for player i, and any $\alpha \in [0, 1]$, let $\sigma_i^{h_0,\alpha}$ be a strategy for player i given by $\sigma_i^{h_0,\alpha}(h) = \alpha$ if $h = h_0$ and $\sigma_i^{h_0,\alpha}(h) = \sigma_i(h)$ otherwise.

Definition 5. Given two strategies $\sigma_i(h)$, $\sigma_i(h')$ for player i, we say $\sigma_i \geq \sigma_i'$ iff $\sigma_i(h) \geq \sigma_i'(h)$ for all histories h.

Definition 6. Given a strategy σ_j for player j and a history h, denote by $V_i(\sigma_j | h)$ the highest continuation utility player i can attain conditional on the history being h and player j using strategy σ_j, i.e.,

$$V_i(\sigma_j | h) = \sup_{\sigma_i} U_i(\sigma_i, \sigma_j | h).$$

Definition 7. Given a strategy σ_j for player j and a history h, denote by $V_i(\sigma_j | h)$ the highest continuation utility player i can attain conditional on the history being h and player j using strategy σ_j, if player i is restricted to not surrendering in the current period, i.e.,

$$\tilde{V}_i(\sigma_j | h) = -c_i(\theta(h)) + \delta E(V_i(\sigma_j | h') | h)$$

Lemma 2. *Let $\sigma_j \geq \sigma'_j$ be two strategies for player j, let σ_i be a strategy for player i and let h be any history. Then*

$$U_i(\sigma_i, \sigma_j|h) \leq U_i(\sigma_i, \sigma'_j|h).$$

Proof. Assume h is a history for time t_0. Then

$$U_i(\sigma_i, \sigma_j|h) - U_i(\sigma_i, \sigma'_j|h) =$$

$$\sum_{t=t_0}^{\infty} \delta^{t-t_0} \int Q(\sigma_i, \sigma'_j, (\theta)^t)(\sigma_j((\theta)^t) - \sigma'_j((\theta)^t))\sigma_i((\theta^t)) \left(U_i\left(\sigma_i^{(\theta)^t,1}, \sigma_j^{(\theta)^t,1}|(\theta)^t\right) - H_1\right) dP((\theta)^t|(\theta)^{t_0}) \leq 0$$

where $Q(\sigma_i, \sigma'_j, (\theta)^t)$ is the probability that the war continues up to time t conditional on the path of the state of the world being $(\theta)^t$ and the players using strategies σ_i, σ'_j respectively. The last inequality follows from the fact that $U_i\left(\sigma_i^{(\theta)^t,1}, \sigma_j^{(\theta)^t,1}|(\theta)^t\right) - H_1 < 0$ by Assumption B5. □

Corollary 1. *Let $\sigma_j \geq \sigma'_j$ be two strategies for player j and let h be any history. Then*

$$V_i(\sigma_j|h) \leq V_i(\sigma'_j, h).$$

Definition 8. We say a strategy σ_i for player i is a *subgame-perfect response* to a strategy σ_j for player j if, for every history h, σ_i is a best response to σ_j in the subgame starting at history h.

Lemma 3. *Let σ_j be a strategy for player j. Then any subgame-perfect response σ_i to σ_j must satisfy: $\sigma_i(h) = 1$ if $\tilde{V}_i(\sigma_j|h) > 0$ and $\sigma_i(h) = 0$ if $\tilde{V}_i(\sigma_j|h) < 0$.*

Proof. If $\tilde{V}_i(\sigma_j|h) > 0$, then $V_i(\sigma_j, h) = \tilde{V}_i(\sigma_j|h) > 0$ and there are strategies σ_i for which $U_i(\sigma_i, \sigma_j|h)$ is arbitrarily close to $V_i(\sigma_j, h)$, hence positive. Then, if σ_i is a subgame-perfect response to σ_j, it must be a best response starting at history h, hence it must attain positive utility starting at history h. This implies $\sigma_i(h) > 0$ and $-c_i(\theta(h)) + \delta E(U_i(\sigma_i, \sigma_j|h')|h) > 0$, whence $\sigma_i(h) = 1$ is optimal.

If $\tilde{V}_i(\sigma_j|h) < 0$, then $V_i(\sigma_j, h) = 0$ and any strategy σ_i with $\sigma_i(h) > 0$ would attain negative continuation utility starting at history h. Hence $\sigma_i(h) = 0$. □

Lemma 4. *Let σ_j, σ'_j be two strategies for player j such that $\sigma_j \geq \sigma'_j$. Let $\sigma_i \in BR_i(\sigma_j)$. Then there is $\sigma'_i \in BR_i(\sigma'_j)$ such that $\sigma'_i \geq \sigma_i$.*

Proof. From Corollary 1 and the definition of $\tilde{V}_i(\sigma_j|h)$, it follows that $\tilde{V}_i(\sigma_j|h) \leq \tilde{V}_i(\sigma'_j|h)$ for all histories h.

Let A_+, A_0 and A_- be the set of histories h for which $\tilde{V}_i(\sigma_j|h) > 0$, $\tilde{V}_i(\sigma_j|h) = 0$ and $\tilde{V}_i(\sigma_j|h) < 0$ respectively, and define A'_+, A'_0 and A'_- analogously for σ'_j. Then $A_+ \subseteq A'_+$ and $A_+ \cup A_0 \subseteq A'_+ \cup A'_0$.

Define σ'_i as follows: $\sigma'_i(h) = 1$ if $h \in A'_+ \cup A'_0$ and $\sigma'_i(h) = 0$ otherwise. Then $\sigma'_i \geq \sigma_i$ by construction, and σ'_i is a best response to σ'_j by Lemma 3. □

Lemma 5. *There are $\underline{M} > -M + \eta$ and $\overline{M} < M - \eta$ such that, in any SPE, player 1 surrenders if $\theta > \overline{M}$ and 2 surrenders if $\theta < \underline{M}$.*

Proof. Assume that $\theta_t \geq M - \eta$, and that player 2 plays a threshold strategy with threshold M_2. As usual, player 1 can guarantee a payoff of 0 by surrendering.

Suppose that player 1 does not surrender immediately. There are two possible types of outcomes. Either player 1 surrenders at some time $t' > t$, or player 2 surrenders at some time $t' \geq t$. In the first case, player 1's continuation utility is strictly negative as she only pays flow costs and receives no prize. In the second case, it must be that $\theta_{t'} \leq M_2$. Player 1's utility is

$$\delta^{t'-t} H_1 - \sum_{s=t}^{s=t'-1} \delta^{s-t} c_1(\theta_s).$$

Recall that, by Assumption A3, $|\theta_{s+1} - \theta_s| \leq \eta$ for all s. Then $M - \eta - M_2 \leq |\theta_{t'} - \theta_t| \leq (t'-t)\eta$. Then

$$\sum_{s=t}^{t'-1} c_1(\theta_s) \geq (t'-t)c_1(M_2) \geq c_1(M_2)\frac{M - M_2 - \eta}{\eta} > H_1,$$

where the last inequality uses Assumption B4. Hence

$$\delta^{t'-t} H_1 < \sum_{s=t}^{t'-1} \delta^{t'-t} c_1(\theta_s) \leq \sum_{s=t}^{t'-1} \delta^{s-t} c_1(\theta_s).$$

Hence player 1's continuation utility is also negative in this case. Thus, player 1 would strictly prefer to surrender if $\theta_t \geq M - \eta$. By continuity, player 1 would also strictly prefer to surrender for all $\theta < M - \eta$ close enough to $M - \eta$. Note that, by Lemma 2, player 1 would also prefer to surrender if player 2 used any other strategy that does not violate Lemma 1.

The argument for player 2 is analogous. □

Lemma 6. *Let $\theta_* \in [-M, M_2]$. If player 2 uses a threshold strategy with threshold θ_*, player 1 has an essentially unique subgame-perfect response, which is also a threshold strategy.*

Analogously, let $\theta^* \in [-M_1, M]$. If player 1 uses a threshold strategy with threshold θ^*, player 2 has an essentially unique subgame-perfect response, which is also a threshold strategy.

Proof. We will prove the first statement; the second one is analogous.

Suppose that player 2 uses a threshold strategy with threshold $\theta_* \in [-M, M_2]$. We will denote this strategy by $\sigma_2^{\theta_*}$.

Let $V_1(\theta)$ be the highest continuation utility player 1 can attain conditional on the current state being θ and player 2 using strategy $\sigma_2^{\theta_*}$, i.e.,

$$V_1(\theta) = \sup_{\sigma_1} U_1(\sigma_1, \sigma_2^{\theta_*} | \theta)$$

Note that V_1 only depends on the current state and not on the history of states of the world, since player 2 is not conditioning on the history.

Next, we prove several properties of $V_1(\theta)$ by a recursive argument.

Claim 1. $V_1(\theta)$ is weakly decreasing in θ.

Proof. Let $V_{10}(\theta)$ be given by $V_{10}(\theta) = H_1$ if $\theta \leq \theta_*$ and $V_{10}(\theta) = 0$ otherwise. Let $\mathcal{L}$ denote the set of Lebesgue-measurable functions from $[-M, M]$ to $[0, H_1]$. Define the operator $W : \mathcal{L} \to \mathcal{L}$ by

$$W(g)(\theta) = \begin{cases} H_1 & \text{if } \theta \leq \theta_* \\ \max\left(-c_1(\theta) + \delta E(g(\theta')|\theta), 0\right) & \text{if } \theta \in (\theta_*, M - \eta) \\ 0 & \text{if } \theta \in [M - \eta, M] \end{cases} \tag{1}$$

where $\theta' - \theta | \theta \sim F_\theta$. For each $k \in \mathbb{N}$, define $V_{1k} = W(V_{1(k-1)})$.

Note that, for all g in the domain of W, $W(g)$ is always in the codomain of W by Assumption B5.

We will now make several observations about W. First, V_1 is a fixed point of W. Indeed, for $\theta \in (\theta_*, M - \eta)$, the statement that $W(V_1)(\theta) = V_1(\theta)$ is just the Bellman equation for V_1. For $\theta \leq \theta_*$, $W(V_1)(\theta) = V_1(\theta) = H_1$ by construction. For $\theta \geq M - \eta$, $W(V_1)(\theta) = V_1(\theta) = 0$ by Lemma 5. Of course, note that $V_1 \in \mathcal{L}$ because $V_1(\theta) \in [0, H_1]$ for all θ by Assumption B5, and V_1 is Lebesgue-measurable since, in fact, it must be continuous on $(\theta_*, M]$ by Assumptions A1 and B2.

Second, W has at most one fixed point by the contraction mapping theorem. Indeed, W is Lipschitz with constant $\delta < 1$ if we endow the space $\mathbb{R}^{[-M,M]}$ with the norm $||\cdot||_\infty$.

Third, W is weakly increasing (i.e., if $g \geq h$ everywhere, $W(g) \geq W(h)$ everywhere).

Fourth, note that $V_{11} \geq V_{10}$ by construction. Then $V_{1(k+1)} \geq V_{1k}$ for all k. Hence, for each θ, the sequence $(V_{1k}(\theta))_k$ is weakly increasing in k. Since it is also bounded, it converges pointwise, and the pointwise limit is a fixed point of W by the monotone convergence theorem. Then, by our previous arguments, V_{1k} converges pointwise to V_1.

Fifth, W preserves decreasing-ness: if g is weakly decreasing in θ, so is $W(g)$. For $\theta \in [\theta_*, M-\eta]$, this follows from Assumptions A2, B1 and B5. For other values of θ, it is obvious. Then, since V_{10} is weakly decreasing in θ, V_{1k} is weakly decreasing in θ for all k, and so is V_1. □

Denote $\tilde{V}_1(\theta) = -c_1(\theta) + \delta E(V_1(\theta')|\theta)$.

Claim 2. $\tilde{V}_1(\theta)$ is strictly decreasing in θ.

Proof. This follows from the facts that $V_1(\theta')$ is weakly decreasing in θ' (Claim 1); θ' is FOSD-increasing in θ by Assumption A2; and $c_1(\theta)$ is strictly increasing in θ by Assumption B1. □

Claim 3. $\tilde{V}_1(\theta)$ and $V_1(\theta)$ are continuous for $\theta \in (\theta_*, M]$.

Proof. $\tilde{V}_1(\theta)$ is continuous in θ for the following reasons: $c_1(\theta)$ is continuous by Assumption B2; V_1 is bounded, as $V_1(\theta) \in [0, H_1]$ for all θ; and h_θ is continuous in θ by Assumption A1.

Recall that, for $\theta \in (\theta_*, M]$, $V_1(\theta) = \max(\tilde{V}_1(\theta), 0)$. Then, since $\tilde{V}_1$ is continuous in θ and the function $\max(\cdot, 0)$ is continuous, $V_1(\theta)$ is continuous in θ for all $\theta \in (\theta_*, M]$. □

Now note that, by Lemma 5, $\tilde{V}_1(\theta) < 0$ for $\theta = M - \eta$, and $\tilde{V}_1(\theta)$ is continuous and strictly decreasing in θ by Claims 2 and 3. Then there are two possibilities. Either there is a unique $\theta^* > \theta_*$ for which $\tilde{V}_1(\theta^*) = 0$ or $\tilde{V}_1(\theta) < 0$ for all $\theta > \theta_*$.

By Lemma 3, in the first case, $\sigma_1^{\theta^*}$ is the essentially unique subgame-perfect response to $\sigma_2^{\theta_*}$.[12] In the second case, the unique best response for player 1 is a threshold strategy with threshold θ_*, such that $\sigma_1(\theta_*) = 1$. □

Proof of Proposition 1. First, we prove that there is an essentially unique equilibrium in threshold strategies. By Lemma 6, if one player is using a threshold strategy, the other player must also be using a threshold strategy, and the latter threshold is uniquely determined as a function of the former. Define then two functions $T_1, T_2 : [-M, M] \to [-M, M]$ as follows: if player i uses threshold θ_i, then player j's optimal threshold as found in Lemma 6 is $T_j(\theta_i)$. An equilibrium in threshold strategies is then given by a threshold θ^* for player 1 such that $T_1(T_2(\theta^*)) = \theta^*$.

We will now show that T_1 is weakly increasing. Let $V_1^{\tilde{\theta}}(\theta)$ and $V_{1k}^{\tilde{\theta}}(\theta)$ for all k be as defined in Lemma 6, conditional on player 2 using threshold $\tilde{\theta}$. Note that, given any two values $\tilde{\theta} > \tilde{\theta}'$, $V_{10}^{\tilde{\theta}} \geq V_{10}^{\tilde{\theta}'}$. Moreover, $W^{\tilde{\theta}}(g) \geq W^{\tilde{\theta}'}(f)$ for any function g, and both operators are weakly increasing. Hence $V_{1k}^{\tilde{\theta}} \geq V_{1k}^{\tilde{\theta}'}$ for all k, so $V_1^{\tilde{\theta}} \geq V_1^{\tilde{\theta}'}$ and $T_1(\tilde{\theta}) \geq T_1(\tilde{\theta}')$, i.e., T_1 is weakly increasing.

[12] It is not unique in the sense that any value $\sigma_1(\theta^*) \in [0, 1]$ is optimal.

Next, we argue that, for any $x > y$ such that $T_1(y) > y$, $T_1(x) - T_1(y) < x - y$. In broad strokes, we will make the following argument. By construction, player 1 is indifferent about continuing when the current state is $T_1(y)$ and player 2 uses threshold y. Suppose now that player 2 switches to using a higher threshold $x > y$, and player 1's optimal response requires her to increase her own threshold exactly as much as player 2 did, i.e., to $z = T_1(y) + x - y$. Then, under the new strategy profile, player 1's utility in state z is lower than her utility in state $T_1(y)$ under the old strategy profile, for two reasons: her flow costs are higher, and the Markov process governing the state is more likely to drift to the right. The same problem arises if $z - T_1(y) > x - y$. Hence player 1's optimal response must involve moving her threshold up by less than $x - y$.

Formally, let t_Δ be the function $t_\Delta(\theta) = \theta - \Delta$. Take $\Delta = \tilde{\theta} - \tilde{\theta}'$. For any function V, denote $\overline{V} = V \circ t_\Delta$. For any operator W, define $\overline{W}$ by $\overline{W}(g) = W(g \circ t_\Delta^{-1}) \circ t_\Delta$.

By construction, $\overline{V}_{1k}^{\tilde{\theta}'} = \overline{W}^{\tilde{\theta}'}(\overline{V}_{1(k-1)}^{\tilde{\theta}'})$ for all k, and $\overline{V}_{10}^{\tilde{\theta}'} = V_{10}^{\tilde{\theta}}$.

The crucial observation now is that, for any weakly decreasing function g, $\overline{W}^{\tilde{\theta}'}(g) \geq W^{\tilde{\theta}}(g)$. Indeed,

$$\overline{W}^{\tilde{\theta}'}(g)(\theta) = \begin{cases} H_1 & \text{if } \theta - \Delta \leq \tilde{\theta}' \Leftrightarrow \theta \leq \tilde{\theta} \\ \max\left(-c_1(\theta - \Delta) + \delta E(g(\theta_{t+1} + \Delta)|\theta_t = \theta - \Delta), 0\right) & \text{if } \theta > \tilde{\theta} \end{cases}$$

Note that $-c_1(\theta - \Delta) > -c_1(\theta)$ by Assumption B1, and $\theta_{t+1} + \Delta = (\theta - \Delta) + X + \Delta$ where X has distribution function $F_{\theta-\Delta}$, which is weakly FOSD'd by F_θ by Assumption A2.

It follows that $\overline{V}_{1k}^{\tilde{\theta}'} \geq V_{1k}^{\tilde{\theta}}$ for all k, and hence $\overline{V}_1^{\tilde{\theta}'} \geq V_1^{\tilde{\theta}}$.

Finally, from Lemma 6, we know that $-c_1(\theta) + \delta E(V_1^{\tilde{\theta}}(\theta')|\theta) = 0$ for $\theta = T_1(\tilde{\theta})$. The above argument implies that $-c_1(\theta - \Delta) + \delta E(f(\theta_{t+1} + \Delta)|\theta_t = \theta - \Delta) > 0$ for the same θ, whence $T_1(\tilde{\theta}') + \Delta > T_1(\tilde{\theta})$. This finishes the argument. On the other hand, if $x > y$ and $T_1(y) = y$, a similar argument implies $T_1(x) = x$, so $T_1(x) - T_1(y) = x - y$.

The analogous results are true of T_2. In addition, it is not possible that $T_2(x) = T_1(x) = x$ for any x. Indeed, if this were the case, by Lemma 6, there would be an equilibrium with thresholds $\theta_* = \theta^* = x$, in which both players have $\tilde{V}_i(x) \leq 0$. But in this case $\tilde{V}_1(x) = -c_1(x) + \delta p H_1$ and $\tilde{V}_2(x) = -c_2(x) + \delta(1-p)H_2$, where p is the probability that $\theta_{t+1} > x$ tomorrow, so it would be implied that $0 \geq \tilde{V}_1(x) + \tilde{V}_2(x) = -c_1(x) - c_2(x) + pH_1 + (1-p)H_2$, which contradicts Assumption B6.

Taken all together, these arguments imply that $T_1 \circ T_2$ has at most one fixed point. Indeed, if $\theta^* \neq \theta^{*'}$ are both fixed points of $T_1 \circ T_2$, we would have that $|T_1(T_2(\theta^*)) - T_1(T_2(\theta^{*'}))| \leq |T_2(\theta^*) - T_2(\theta^{*'})| \leq |\theta^* - \theta^{*'}|$ with at least one strict inequality, a contradiction.

Next, we show that $T_1 \circ T_2$ has a fixed point as follows. Take $\overline{\theta}_0^* = M$ and $\overline{\theta}_n^* =$

$T_1(T_2(\overline{\theta}^*_{n-1}))$ for all $n \geq 1$. Clearly $\overline{\theta}^*_0 \geq \overline{\theta}^*_1$. Since $T_1 \circ T_2$ is weakly increasing, it follows that $\overline{\theta}^*_1 \geq \overline{\theta}^*_2 \geq \ldots$. Since the sequence is bounded it must converge to a limit $\overline{\theta}^*$. As we have shown that $T_1 \circ T_2$ is Lipschitz, and hence continuous, it follows that $\overline{\theta}^*$ is a fixed point of $T_1 \circ T_2$.

Finally we rule out other equilibria that are not in threshold strategies. We use a standard argument from supermodular games similar to Milgrom and Roberts [[[cite]]]. Following the notation of the previous paragraph, denote $\overline{\theta}_{*n} = T_2(\overline{\theta}^*_n)$ for all n and $\overline{\theta}_* = T_2(\overline{\theta}^*)$. Also, let $\underline{\theta}_{*0} = -M$, $\underline{\theta}_{*n} = T_2(T_1(\underline{\theta}_{*(n-1)}))$, $\underline{\theta}^*_n = T_1(\underline{\theta}_{*n})$ and denote the limits by $\underline{\theta}_*, \underline{\theta}^*$ respectively.

Since T_1, T_2 are weakly increasing, we have

$$\overline{\theta}^*_0 \geq \ldots \geq \overline{\theta}^* \geq \underline{\theta}^* \geq \ldots \geq \underline{\theta}^*_0$$
$$\overline{\theta}_{*0} \geq \ldots \geq \overline{\theta}_* \geq \underline{\theta}_* \geq \ldots \geq \underline{\theta}_{*0}$$

By Lemma **??**, whenever i plays a strategy higher than σ_i, any best response by j must be weakly lower than j's best response to σ_i. Hence, any strategy played by 1 must be bounded between $\sigma_1^{\underline{\theta}^*_0}$ and $\sigma_1^{\overline{\theta}^*_0}$; any strategy played by 2 must be bounded between $\sigma_2^{\overline{\theta}_{*0}}$ and $\sigma_2^{\underline{\theta}_{*0}}$; and so on.

By induction, any strategy used by 1 must be between $\sigma_1^{\underline{\theta}^*}$ and $\sigma_1^{\overline{\theta}^*}$. But since $\underline{\theta}^* = \overline{\theta}^* = \theta^*$, there is a (essentially) unique equilibrium strategy for player 1. The same argument applies to player 2. □

Proof of Proposition 2. For (i), take two cost functions c_1, $\hat{c}_1$ for player 1 such that $\hat{c}_1(\theta) < c_1(\theta)$ for all θ. (The cases where H_1 increases or c_2 or H_2 change are analogous.) Assume that player 2 is playing a threshold strategy with threshold θ_*. Using the notation developed in Proposition 1, let $V_1(\theta)$ and $\hat{V}_1(\theta)$ be the value functions for player 1 when her cost function is $c_1(\theta)$ and $\hat{c}_1(\theta)$, respectively. We will similarly refer to the analogues of W, T_1 under the cost function $\hat{c}_1$ as $\hat{W}$, $\hat{T}_1$, respectively.

Note that $\hat{W}(g) \geq W(g)$ for all $g \in \mathcal{L}$. Hence $\hat{W}(V_1) \geq W(V_1) = V_1$. As argued in Proposition 1, $\hat{W}$ is increasing and $\hat{V}_1$ must be the limit of $\hat{W}^k(g)$ for any g by the Contraction Mapping Theorem. Hence

$$V_1 \leq \hat{W}(V_1) \leq \hat{W}^2(V_1) \leq \ldots \nearrow \hat{V}_1,$$

whence $\hat{V}_1 \geq V_1$. From this and the fact that $\hat{c}_1(\theta) < c_1(\theta)$ for all θ it follows that $\hat{\tilde{V}}_1(\theta) > \tilde{V}_1(\theta)$ for all θ. Assume that $T_1(\theta_*) > \theta_*$. Then we have $\hat{V}_1(\theta) > V_1(\theta)$ for all $\theta \in (\theta_*, T_1(\theta_*))$. By the continuity of $\hat{V}_1$, $\hat{V}_1(\theta) > V_1(\theta) \geq 0$ for all θ in a neighborhood of $T_1(\theta_*)$ as well, so $\hat{T}_1(\theta_*) > T_1(\theta_*)$.

Let θ_*, θ^* denote the equilibrium thresholds when player 1's cost function is c_1, and let $\hat{\theta}_*$, $\hat{\theta}^*$ denote the equilibrium thresholds when player 1's cost function is $\hat{c}_1$. Since nothing about player 2's problem has changed, T_2 remains unchanged. $\hat{\theta}_*$, $\hat{\theta}^*$ are characterized by the conditions that $\hat{\theta}^*$ be a fixed point of $\hat{T}_1 \circ T_2$ and $\hat{\theta}_* = T_2(\hat{\theta}^*$. As $\theta^* = T_1(\theta_*) > \theta_*$ by Proposition 1, we have $\hat{T}_1(\theta_*) > \theta^*$.

Because T_1 and T_2 are weakly increasing, we have that

$$\theta^* < \left(\hat{T}_1 \circ T_2\right)(\theta^*) \leq \left(\hat{T}_1 \circ T_2\right)^2 (\theta^*) \leq \ldots \nearrow \hat{\theta}^*$$

Hence $\hat{\theta}^* > \theta^*$. By an analogous argument $\hat{\theta}_* > \theta_*$.

As for the claim that $\hat{\theta}^* - \hat{\theta}_* > \theta^* - \theta_*$, recall that, in Proposition 1, we argued that $T_i(x) - T_i(y) < x - y$ whenever $x > y$ are such that $T_i(y) > y$. Here, that implies

$$\hat{\theta}_* - \theta_* = T_2(\hat{\theta}^*) - T_2(\theta^*) < \hat{\theta}^* - \theta^*,$$

which yields the result.

The proof of (ii) is similar to (i). Briefly, denoting by $(\hat{f}_\theta)_\theta$ a new set of transition probabilities, and by $\hat{W}_i$, $\hat{V}_i$ and $\hat{T}_i$ the new operators, value functions and threshold mappings under the new transition probabilities, we can show that $\hat{W}_1(g) \leq W(g)$ for any weakly decreasing g, and $\hat{W}_2(g) \geq W_2(g)$ for any weakly increasing g. Hence $\hat{V}_1 \leq V_1$ and $\hat{V}_2 \geq V_2$, for fixed conjectures about the other player's behavior, from which it follows that $\hat{T}_1 \leq T_1$ and $\hat{T}_2 \leq T_2$. By a similar argument as above, this implies $\hat{\theta}_* \leq \theta_*$ and $\hat{\theta}^* \leq \theta^*$.

(iii) Let $\hat{H}_1 = \rho H_1$, $\hat{H}_2 = \rho H_2$ with $\rho > 1$. Note that part (i) already implies $\hat{\theta}^* - \hat{\theta}_* > \theta^* - \theta_*$.

Suppose the claim is not true, and WLOG suppose further that $\hat{\theta}_* \geq \theta_*$.

I think it's easider not to go to equilibrium thresholds, and instead do the following. Let $\theta_1^* = \hat{T}_1(\theta_*)$ and $\theta_{*1} = \hat{T}_2(\theta_1^*)$. Then we can show that, if $\theta_{*1} \geq \theta_*$, then the (contradictory) claim holds, and if $<$, it does not. So suppose it holds.

Try to prove it using the derivative. i.e., study the change in T_1 as H_1 moves up by ϵ, and possibly as c_1 changes by λd_1.

Marginal change of T_1 depends on 3 things: -impact of H_1 in utility function (probably a function of delay until H_1 is hit, which will be the same for both players *if we assume f symmetric around 0*), which affects the partial of V_1 wrt H_1 everywhere -slope of V_1 at $T_1(\theta^*)$: higher slope means lower change in threshold from increasing V_1 everywhere -impact of c_1 in utility function

So, even in the absence of asymmetries in cost changes, the change in V_i as H is varied will be the SAME for both players. So if V_i' at the threshold is equal for both players, or higher for player 1, we get the result for free. If the derivative is higher for player 2, then we need to figure out the impact of the cost function.

[(iii)] Suppose $(F_\theta)_\theta$ is such that $\mu(\theta) = 0$ for all θ. Then, if H_1, H_2 increase proportionally, θ_* decreases and θ^* increases.

□

Proof of Proposition 3. For the first part, note that, for any $t > 0$, and any $\kappa > 0$ such that t is a multiple of κ, $V_i(\theta)$ must satisfy the time-t Bellman equation approximately (approximately in the sense that it will not hold exactly when the game stops at a time that is not a multiple of t, but it should be a small error). As $\kappa \to 0$ each of these conditions should give us a uniform bound on the limit of V_i, and hence on the limit of T_i.

For the second part, note that both $V_i(\theta_t)$, i's equilibrium continuation utility starting in state θ_t, and $P_i(\theta_t)$, i's equilibrium probability of winning starting in state θ, are themselves drift-diffusion processes by Îto's lemma:

$$dV_i(\theta_t) = \left(\mu(\theta_t)V_i'(\theta_t) + \frac{\sigma^2(\theta_t)}{2}V_i''(\theta_t)\right)\Big(dt + \sigma(\theta_t)V_i(\theta_t)dB_t \tag{2}$$

$$dP_i(\theta_t) = \left(\mu(\theta_t)P_i'(\theta_t) + \frac{\sigma^2(\theta_t)}{2}P_i''(\theta_t)\right)\Big(dt + \sigma(\theta_t)P_i(\theta_t)dB_t \tag{3}$$

At the same time, it follows from the continuous-time Bellman equation for V_i that

$$0 = E_t(E_{t'}(V_i(\theta))) \approx (t' - t)(-c_i(\theta) + e^{-\gamma(t'-t)}\left(V_i(\theta) + (t' - t)E(dV_i(\theta))\right)$$
$$0 = -c_i(\theta) - \gamma V_i(\theta) + E(dV_i(\theta))$$

and it follows from the law of iterated expectations that $E(dP_i(\theta)) = 0$.

Taking expectation of Equations 2 and 3 conditional on the value of θ_t, we have

$$c_i(\theta) + \gamma V_i(\theta) = E(dV_i(\theta)) = \mu(\theta)V_i'(\theta) + \frac{\sigma^2(\theta)}{2}V_i''(\theta)$$
$$0 = dP_i(\theta) = \mu(\theta)P_i'(\theta) + \frac{\sigma^2(\theta)}{2}P_i''(\theta)$$

Let $\theta_*(\kappa)$, $\theta^*(\kappa)$ be the equilibrium thresholds from Proposition **??** as a function of κ. Then there are θ_{*0}, θ^{*0} such that $\theta_*(\kappa) \to \theta_{*0}$ and $\theta^*(\kappa) \to \theta^{*0}$ as $\kappa \to 0$.

In addition, each player's expected utility $U_i(\theta)$ and probability of winning $P_i(\theta)$, starting at a given θ, can be calculated in $[\theta_{*0}, \theta^{*0}]$ as the solution to the following ODEs:

$$\mathrm{E}(\mathrm{dV}_i(\theta)) =$$

$$\mathrm{U"}_i(\theta) = \frac{2}{\mu(\theta)^2+\sigma^2}\left(c_i(\theta) - \mu(\theta)U_i'(\theta)\right)$$

$$P_i''(\theta) = -\frac{2}{\mu(\theta)^2+\sigma^2}\mu(\theta)P_i'(\theta)$$

given the boundary conditions $U_1(\theta_{*0}) = H_1$; $U_2(\theta_{*0}) = 0$; $U_1(\theta^{*0}) = 0$; $U_2(\theta^{*0}) = H_2$; $P_1(\theta_{*0}) = 1$; $P_2(\theta_{*0}) = 0$; $P_1(\theta^{*0}) = 0$; $P_2(\theta^{*0}) = 1$.

In particular, when $\mu \equiv 0$, the solution in $[\theta_{*0}, \theta^{*0}]$ reduces to

$$U_1(\theta) = \frac{2}{\sigma^2}\int_\theta^{\theta^{*0}}\left(\int_\lambda^{\theta^{*0}} c_1(\omega)d\omega\right)d\lambda = \frac{2}{\sigma^2}\int_\theta^{\theta^{*0}}(\lambda-\theta)c_1(\lambda)d\lambda \qquad P_1(\theta) = \frac{\theta^{*0}-\theta}{\theta^{*0}-\theta_{*0}}$$

$$U_2(\theta) = \frac{2}{\sigma^2}\int_{\theta_{*0}}^{\theta}\left(\int_{\theta_{*0}}^{\lambda} c_2(\omega)d\omega\right)d\lambda = \frac{2}{\sigma^2}\int_{\theta_{*0}}^{\theta}(\theta-\lambda)c_2(\lambda)d\lambda \qquad P_2(\theta) = \frac{\theta-\theta_{*0}}{\theta^{*0}-\theta_{*0}},$$

and the thresholds θ_{*0}, θ^{*0} are determined by the conditions $U_1(\theta_{*0}) = H_1$, $U_2(\theta^{*0}) = H_2$.

□

References

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

Burdzy, Krzysztof, David M Frankel, and Ady Pauzner, "Fast equilibrium selection by rational players living in a changing world," *Econometrica*, 2001, *69* (1), 163–189.

Chatterjee, Kalyan and Larry Samuelson, "Bargaining with Two-Sided Incomplete Information: An Infinite Horizon Model with Alternating Offers," *The Review of Economic Studies*, 1987, *54* (2), 175–192.

Egorov, Georgy and Bård Harstad, "Private Politics and Public Regulation," *The Review of Economic Studies*, 2017, p. rdx009.

Fudenberg, Drew and Jean Tirole, "A Theory of Exit in Duopoly," *Econometrica*, 1986, pp. 943–960.

Kreps, David M and Robert Wilson, "Reputation and Imperfect Information," *Journal of Economic Theory*, 1982, *27* (2), 253–279.

Krishna, Vijay and John Morgan, "An Analysis of the War of Attrition and the All-Pay Auction," *Journal of Economic Theory*, 1997, *72* (2), 343–362.

Milgrom, Paul and John Roberts, "Predation, Reputation, and Entry Deterrence," *Journal of Economic Theory*, 1982, *27* (2), 280–312.

_ **and** _ , "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica: Journal of the Econometric Society*, 1990, pp. 1255–1277.

Ortner, Juan, "Optimism, delay and (in)efficiency in a stochastic model of bargaining," *Games and Economic Behavior*, 2013, *77* (1), 352–366.

_ , "A Continuous Time Model of Bilateral Bargaining," 2016.

_ , "A theory of political gridlock," *Theoretical Economics*, 2017, *12* (2), 555–586.

Ponsati, Clara and József Sákovics, "The war of attrition with incomplete information," *Mathematical Social Sciences*, 1995, *29* (3), 239–254.

Slantchev, Branislav L, "The Principle of Convergence in Wartime Negotiations," *American Political Science Review*, 2003, *97* (4), 621–632.

Smith, J Maynard, "The Theory of Games and the Evolution of Animal Conflicts," *Journal of Theoretical Biology*, 1974, *47* (1), 209–221.

Topkis, Donald M, "Equilibrium Points in Nonzero-Sum n-Person Submodular Games," *Siam Journal of Control and Optimization*, 1979, *17* (6), 773–787.