Evolving Wars of Attrition*

Germán Gieczewski†

June 2024

Abstract

This paper models a war of attrition that evolves over time. Two players fight over a prize until one surrenders. The flow costs of fighting depend on a state variable that is public but changes stochastically as the war unfolds. In the unique equilibrium, each player surrenders when the state becomes adverse enough; for intermediate states, both players fight on. In an extension, the baseline model is augmented to allow the players to unilaterally concede part of the prize. Such concessions can be beneficial if they disproportionately sap the opponent's incentive to fight. The evolving war of attrition with concessions yields predictions regarding delay and the eventual division of the prize that differ from conventional models of bargaining as well as reputational wars of attrition.

JEL codes: C73, C78

Keywords: war of attrition, concessions, equilibrium selection, evolving state

*I am thankful to Avidit Acharya, Jonathan Bendor, Peter Bils, Robert Gulotty, Matías Iaryczower, Gleason Judd, Adam Meirowitz, Wolfgang Pesendorfer, Kristopher Ramsay, Juuso Toikka, seminar participants at Stanford, NYU and Vanderbilt and conference participants at APSA, MPSA and Stony Brook for helpful comments and suggestions.

†Department of Politics, Princeton University. Email: `germang@princeton.edu`. Address: 311 Fisher Hall, Princeton NJ 08540.

1 Introduction

The war of attrition is a model with wide applicability in economics and politics. Two firms in a duopoly engaging in a price war (Fudenberg and Tirole, 1986); an activist group boycotting a firm (Egorov and Harstad, 2017); a government facing a protest movement; two political parties delaying a fiscal adjustment in hopes of shifting blame onto the other party (Alesina and Drazen, 1991); or a protracted military conflict, such as the Western Front of World War I, can all be understood as wars of attrition.

In the classic war of attrition (Smith, 1974; Hendricks, Weiss and Wilson, 1988), two players fight over a prize by paying flow costs until one surrenders. Both the value of the prize and the costs of continuing the war are deterministic and publicly known. As is well known, this model allows for equilibria in which either player surrenders immediately, as well as an equilibrium in which either player may win, and the length of the war is random (Fudenberg and Tirole, 1991). Therefore, the classic model is silent on the crucial questions of which player should win the war, and how long it will take before the loser concedes.[1]

The assumption that payoffs are perfectly predictable not only leads to a multiplicity of equilibria. It also elides a crucial feature of most practical applications: the fact that, as they unfold, wars change in unexpected ways that may favor either player. For example, take the case of two ride-sharing companies that cut prices in an attempt to drive each other out of business. Both firms expect to lose money, but cannot know how their market shares will change over time, or how much capital they can raise to survive; this information is only revealed as the price war progresses. Similarly, in the middle of a prolonged protest, neither the government nor the protesters know how public opinion will shift as the stalemate drags on. And weather changes (Winters, 2001), as well as unexpected battlefield outcomes, can change the course of (literal) wars.

Motivated by these observations, this paper models a war of attrition that evolves over time. The model presented is identical to the classic war of attrition, except that there is a state of the world, θ_t, which is commonly observed at all times, changes stochastically over time, and affects the flow costs that players must pay to continue the war. In the baseline model, θ_t parameterizes the extent to which current conditions in the war favor one player over the other: high values of θ_t mean high costs for player 1 and low costs for player 2, while low values of θ_t mean the opposite.

[1] Some papers in this literature (e.g., Smith (1974), Bliss and Nalebuff (1984), Maskin and Tirole (1988), Kapur (1995), Wang (2009), Pitchford and Wright (2012), Montez (2013)) sidestep the issue of equilibrium multiplicity by focusing on the mixed strategy equilibrium with no instantaneous concession. This is a sensible approach only when the players are symmetric; in the asymmetric case, the mixed equilibrium gives the weaker player (that is, the one with higher cost-to-prize ratio) a higher chance of winning—an implausible prediction.

Under mild assumptions, the game has a unique equilibrium, with the following structure. So long as θ_t lies in a certain interval, referred to as the *disputed region*, both players continue fighting, and they have strict incentives to do so. When θ_t reaches an extreme enough value as to leave the disputed region, the disadvantaged player surrenders. Both players have positive winning probabilities, yet the equilibrium is in pure strategies, and there is no immediate surrender—in fact, under some conditions, the length of the war is bounded away from zero. In these ways, the equilibrium differs qualitatively from those obtained in the classic war of attrition as well as reputational perturbations of it (Abreu and Gul, 2000). The equilibrium also has sensible comparative statics: if a player's prize increases or her cost decreases, her probability of winning increases, both because she is more willing to fight and because her increased belligerence makes victory harder to attain for the opponent.

The baseline model is closely related to the "war of information" modeled in Gul and Pesendorfer (2012). Gul and Pesendorfer model two parties that provide costly information to a voter about which party is better; the voter's changing posterior serves as the state variable. In equilibrium, only the "trailing" party wants to provide information, which leads to a specific formulation of the cost functions. The model in this paper extends Gul and Pesendorfer (2012) by allowing simultaneously for general cost functions, discounting, and a state variable that may drift in one player's favor or be multi-dimensional (see Section 5). In particular, I provide conditions on these underlying objects that are general enough as to be tight, in a sense I make precise.

I then use the baseline model to perform two exercises. First, I consider the limit of the solution as the movement of the state θ_t becomes arbitrarily slow. This limit equilibrium is an equilibrium of the classic war of attrition, albeit one augmented with a payoff-irrelevant, changing state variable. I show that, if the players' cost-prize ratios differ, the stronger player (with a higher prize or lower cost) wins immediately.[2] However, when the players are evenly matched, the limit equilibrium leverages the state variable θ_t as a coordination device, and it features less average delay than the mixed equilibrium of the classic war of attrition, though still a positive amount.

Second, I show how the evolving war of attrition can be extended to allow for additional actions besides continuing to fight and surrendering completely. In particular, I allow the players to unilaterally concede part of the prize to the opponent, and then keep fighting over the rest. For example, in the context of a protest, the government may cede to some, but not all, of the protesters' demands, in an attempt to defuse the protest at the lowest possible cost. This exercise is uniquely tractable in my setting, relative to reputational models, because a

[2] This matches the predictions of reputational models (e.g., Abreu and Gul (2000)), as long as the players' probabilities of being "commitment types" are taken to zero at similar speeds.

concession does not carry signaling content.[3] I characterize the optimal use of concessions when only one player is able to make concessions, and provide partial results for the case in which both can do so. The general principle that arises is that partial concessions can be worthwhile, but only if they reduce the opponent's incentive to fight strictly more than the conceding player's. In particular, making unilateral concessions is never useful if both players value different parts of the prize equally (for example, if the prize is simply a pot of money), but it may be optimal if the prize is heterogeneous, with the two players valuing certain parts disproportionately (as may be the case with territory). Furthermore, it may be the case that in equilibrium some concessions are made, but they do not exhaust the prize, so a smaller conflict over the remainder follows. When both players can make concessions, the setup I study can be taken as a model of bargaining under extreme lack of commitment, meaning that a concession by one player cannot be conditioned on the opponent giving up something else in return.

Besides the works already mentioned, this paper is related to four broad strands of literature. First, many variants of the war of attrition obtain equilibrium selection by adding reputational concerns. This approach, applied to exit in duopoly (Fudenberg and Tirole, 1986), entry deterrence (Kreps and Wilson, 1982; Milgrom and Roberts, 1982) and bargaining (Abreu and Gul, 2000), generally yields a unique equilibrium when the players have a positive probability of being "irrational" types who never surrender. A central result of this literature is that, even when extreme types are rare, the incentive to *pretend* to be extreme shapes the equilibrium behavior of all types.[4] In our model, a related, but less influential role, is played by dominance regions at very high or low values of θ_t.

Whereas in reputational models there is private information, and fighting is a costly signal of resolve, in my model there is symmetric uncertainty, and fighting is a gamble that the war will turn in the player's favor. This distinction is empirically relevant. For instance, in a duopoly, a credible revelation of financial statements, or an act of corporate espionage, could radically alter a reputation-driven price war or end it immediately, while it would have no impact on a war sustained by shared expectations of an uncertain future. As noted above, the two settings diverge further if additional actions are available: a partial concession may be a smart play that saps the opponent's motivation to fight, but in a reputational model, it may be taken to signal weakness.

Secondly, there is a small but growing literature on dynamic games with a changing state of the world. For instance, Ortner (2016) studies a game of bargaining with alternating offers

[3] By contrast, adding partial concessions to the classic war of attrition would not yield sharp predictions due to the extreme multiplicity of equilibria.

[4] This striking result has some extreme implications. For instance, if one player's probability of being irrational is arbitrarily small, and the other's is zero, the second player must surrender immediately.

where the bargaining protocol (i.e., the identity of the player making offers) is driven by a Brownian motion. Ortner (2017) considers two political parties bargaining in the shadow of an election; if an agreement is reached, the result affects the relative popularity of the parties, which otherwise evolves as a Brownian motion and eventually determines the outcome of the election. In Ortner (2013), optimistic players bargain over a prize whose value changes over time. The closest paper in this group is a contemporaneous paper by Georgiadis, Kim and Kwon (2022), which also models a war of attrition with a changing state of the world that affects both players. The crucial difference is that, in Georgiadis et al. (2022), changes in the state affect both players *equally*, while in this paper they affect the players in *opposite* ways. An interpretation in the duopoly setting is that the state in Georgiadis et al. (2022) tracks the size of the market, while in this paper it tracks market shares. This difference leads to contrasting results: in Georgiadis et al. (2022), the equilibrium is unique only if the players' incentives to fight differ substantially, and in every equilibrium the identity of the first player to quit is known in advance. In Section 5, I show that both their model and my benchmark model are nested in a model with a two-dimensional state with a common-values dimension and an adversarial dimension. That model reduces to my benchmark model when only the second dimension changes over time, and to Georgiadis et al. (2022) when only the first one changes. I find that, in the general case in which both dimensions change over time, my results survive—that is, there is a unique equilibrium, and either player may be the first to quit.

Another cluster of papers in this vein studies *races* in which two players exert effort in order to move the state in their favor. In some of these models, each player is at a certain distance from a "finish line" (Harris and Vickers, 1985). More closely related are *tug-of-war* models in which effort pushes a common state towards a player's preferred finish line (Harris and Vickers, 1987; Budd, Harris and Vickers, 1993; Moscarini and Smith, 2011; Cao, 2014). These models are connected to this paper both by motivating examples (e.g., a price war) and the gist of their results: the fact that a disfavored enough player quits in our model is related a common result in this literature, whereby the laggard in a race tends to exert lower effort due to a discouragement effect. However, these models differ from ours in that the effort choice is more flexible (usually continuous) than simply continuing or quitting, and equilibria in these models are not unique under general conditions.[5]

Third, the paper is related to a body of work on conflict in international relations (Smith, 1998; Slantchev, 2003a,b; Powell, 2004). The closest paper in this family, Smith (1998),

[5]For instance, Harris and Vickers (1987) and Moscarini and Smith (2011) give conditions for uniqueness of a symmetric equilibrium in a tug-of-war, but asymmetric equilibria also exist. Budd et al. (1993) shows uniqueness only when the players are very impatient or the state evolves very noisily, while Cao (2014) assumes quadratic costs and no discounting.

models a war in which the players fight over a sequence of "forts", the players' payoffs depend on how many forts they hold, and either player may surrender at any time. While this setting is *prima facie* similar to the one in this paper, the model in Smith (1998) has a potentially large set of equilibria that cannot be fully characterized in general.

Finally, the present paper also contributes to the literature on supermodular games started by Topkis (1979) and Milgrom and Roberts (1990). My approach exploits the fact that the war of attrition is a supermodular game when players' strategies are ordered in opposite ways (i.e., player 1's "high" strategy is to continue while player 2's is to surrender). The way in which perturbing the state leads to equilibrium uniqueness is also reminiscent of results in global games (Morris and Shin, 1998) and related work (Burdzy, Frankel and Pauzner, 2001).

The paper proceeds as follows. Section 2 presents the baseline model in discrete time, characterizes its equilibrium and comparative statics, and then presents a continuous time version. Section 3 discusses the results in the context of existing models, and analyzes the limit of the solution as the state is made to change arbitrarily slowly. Section 4 extends the model to allow for (one-sided) partial concessions. Section 5 presents the model with a multi-dimensional state. Section 6 concludes. All the proofs are found in Appendices A and B. Appendix C gives an example of an equilibrium with two-sided partial concessions.

2 The Model

There are two players, 1 and 2. We first set up a discrete-time model with infinite horizon: $t = 0, 1, \ldots$ before taking the limit to continuous time.[6]

In each period, each player can choose to continue (C) or surrender (S). There is a state of the world $\theta_t \in [-M, M]$ which is common knowledge at all times. θ_t represents how favorable the current conditions are to either player: a high θ_t favors 2, while a low θ_t favors 1. The initial θ_0 is given by Nature. Then it evolves according to a Markov process described by

$$P(\theta_{t+1} - \theta_t \leq x | \theta_t) = F_{\theta_t}(x),$$

where $F_\theta : \mathbb{R} \to [0, 1]$ is an absolutely continuous c.d.f. with corresponding density f_θ.

We assume that, for some value of $\eta > 0$:

A1 f_θ is continuous in θ for all $\theta \in [-M, M]$. More precisely, the mapping $\theta \mapsto f_\theta$ is continuous, taking the 1-norm in the codomain.

[6] The continuous time setting is more analytically tractable, but studying limits of discrete time equilibria rather than working directly in continuous time helps to sidestep some technical issues.

A2 F_θ is weakly FOSD-increasing in θ for all $\theta \in [-M+\eta, M-\eta]$. In other words, if $\theta, \theta' \in [-M+\eta, M-\eta]$ and $\theta \geq \theta'$ then $F_\theta(x) \leq F_{\theta'}(x)$ for all x.

A3 For all $\theta \in [-M, M]$, $\text{supp}(F_\theta)$ is a closed interval with nonempty interior, contained in $[-\eta, \eta] \cap [-M-\theta, M-\theta]$.

A1 says that small changes in θ_t induce small changes in θ_{t+1}, while A3 assumes that θ_t does not change by more than $\pm\eta$ between periods, and always stays within $[-M, M]$. The only substantively important assumption is A2, which rules out mean-reverting processes. (It does, of course, allow for random walks.) Substantively, A2 fits settings where advantages are persistent or even tend to snowball, as might be the case in wars, contests, etc., but it excludes processes with transient or cyclical changes.[7] Technically, the upshot of A2 is that the prospect of θ_t going to extreme values has strategic bite. In contrast, if A2 were severely violated (i.e., if F_θ were highly mean-reverting), even a substantial swing in θ_t may not affect the strategic calculus much, as the players would expect the state to quickly "return to normal". However, our main results can hold even if A2 is slightly violated, as discussed below.

Payoffs are as follows. In any period in which the war is ongoing, each player i pays a flow cost $c_i(\theta)$. If either player chooses to surrender at time t, players do not pay flow costs in that period and the war ends. When the war ends, the winner receives an instantaneous payoff H_i and the loser receives $L_i < H_i$, normalized to 0.[8] If both players surrender on the same turn, they both lose.[9] The players have a common discount factor $\delta \in [0, 1]$.[10] Hence, i's lifetime payoff if the war ends at time T is:

$$U_i(\sigma_i, \sigma_j) = \delta^T H_i \mathbb{1}_{\{i \text{ wins}\}} - \sum_{t=0}^{T-1} \delta^t c_i(\theta_t).$$

We assume the following about the players' payoff functions:

B1 $c_1(\theta)$ is strictly increasing in θ, and $c_2(\theta)$ is strictly decreasing in θ.

B2 $c_1(\theta)$, $c_2(\theta)$ are C^1.

B3 There are M_1, M_2 such that $-M < -M_1 < 0 < M_2 < M$ and $c_1(-M_1) = c_2(M_2) = 0$.

[7] However, in the multi-dimensional model presented in Section 5, the evolution of the state needs to be mean-reinforcing only in one dimension; the rest can accommodate more general behavior.

[8] A player who receives H_i from winning, L_i from surrendering and a flow payoff $-c_i(\theta)$ while the war continues has identical incentives to one who receives $H_i + \rho$ from winning, $L_i + \rho$ from losing and flow payoff $-c_i(\theta) + \rho(1-\delta)$ while the war continues.

[9] This assumption simplifies some formal arguments but is not essential, as players never surrender simultaneously in equilibrium.

[10] $\delta = 1$ does not lead to issues involving infinite utility since it is never rational to continue the war forever.

B4 The following inequalities are satisfied:

$$c_2(-M_1)\frac{M - M_1 - \eta}{\eta} > H_2$$
$$c_1(M_2)\frac{M - M_2 - \eta}{\eta} > H_1.$$

B5 $H_i > \frac{-c_i(\theta)}{1-\delta}$ for $i = 1, 2$ and all $\theta \in [-M, M]$.

B6 $\delta H_i > c_1(\theta) + c_2(\theta)$ for $i = 1, 2$ and all $\theta \in [-M, M]$.

Substantively, Assumption B1 says that player 1 is favored when θ is low, and vice versa. Thus, changes in the state affect the players in opposite ways. (Appendix **??** generalizes the results to a setting with a multidimensional state that can affect payoffs in richer ways.) B3 says that, for favorable enough values of θ, players actively enjoy fighting. B4 guarantees that, when the state is so favorable to one player that she will be in favor of fighting for a long time, it is best for the other player to surrender. B5 says that players never enjoy fighting so much that they would rather continue fighting than win immediately. Finally, B6 guarantees that flow costs are low enough that at least one player is always willing to continue if the war is expected to end in the next period.

Out of these assumptions, B3 is the least innocuous and hence worth discussing. In some contexts, it is reasonable that a player facing a favorable state of the world would obtain positive net flow payoffs from continuing the war, relative to surrendering. For example, in a price war between two duopolists, Assumption B3 reflects that, when a firm has high enough market share, it can turn a profit even before the other firm exits the market. In a dispute between a firm and an activist group, in which θ represents the state of public opinion, the firm's profits might increase if public opinion turns against the activist group, in which case the boycott turns into free publicity for the firm; conversely, if public opinion turns against the firm, activists may obtain (possibly non-pecuniary) payoffs from hurting the firm's bottom line, or from increased funding or exposure for their other causes.

In other examples—for instance, if the game represents the siege of a city by an attacking army—Assumption B3 is less plausible. However, we would obtain much the same results by assuming that, if the state of the world becomes extreme enough—formally, if θ_t goes above (below) some threshold—then the war is over in a material sense and player 1 (2) is forced to surrender, for instance, due to bankruptcy or death.

Our equilibrium concept is Subgame Perfect Equilibrium (SPE). In general, we denote a strategy for player i by ψ_i, where $\psi_i(h)$ is the probability that i continues at history h. If ψ_1 is such that player 1 continues iff $\theta(h) \le \theta^*$, we say ψ_1 is a threshold strategy with threshold

θ^*. Similarly, if ψ_2 is such that player 2 continues iff $\theta(h) \geq \theta_*$, ψ_2 is a threshold strategy with threshold θ_*.

Analysis

Our first main result characterizes the unique subgame perfect equilibrium (SPE) of the game.

Proposition 1. *There is an essentially*[11] *unique SPE. The equilibrium is in threshold strategies: there are $\theta_* < \theta^*$ such that player 1 surrenders whenever $\theta_t > \theta^*$, player 2 surrenders whenever $\theta_t < \theta_*$, and neither player surrenders when $\theta_t \in (\theta_*, \theta^*)$.*

The equilibrium partitions the set of possible states $[-M, M]$ into three intervals: player 2's *surrender region*, $[-M, \theta_*)$; player 1's surrender region, $(\theta^*, M]$; and between them the *disputed region*, $[\theta_*, \theta^*]$, in which both players choose to continue the war.

The intuition behind the proof of Proposition 1 is as follows. First, if θ is very low, then player 1 never has an incentive to surrender: even if she expects to eventually lose, she should stay in the war until θ goes above $-M_1$ (i.e., until her costs become positive). Then, since the gap between $-M_1$ and $-M$ is large (Assumption B4), there will be low enough values of θ for which player 2 is forced to surrender immediately, since waiting for player 1 to surrender will be too costly. Similarly, for very high values θ, player 2 will never want to surrender, so player 1 must concede immediately.

For values of θ between these extremes, behavior will depend on expectations about the other player's strategy. For instance, if 2 plays a "hawkish" strategy in equilibrium—that is, she surrenders only for a small set of values of θ—this incentivizes 1 to play a "dovish" strategy, which surrenders at a large set of values of θ, and vice versa. Formally, the game is supermodular, if we order the strategy sets so that surrender is the "high" action for player 1 and the "low" action for player 2. Hence, as is standard in supermodular games, there is a greatest and a smallest equilibrium that all other equilibria are bounded between (Milgrom and Roberts, 1990). Because the best response to a threshold strategy is another threshold strategy, and the greatest and the smallest equilibrium can be obtained as iterated best responses to threshold strategies, they are themselves in threshold strategies.

The final step of the proof is to show that the greatest and the smallest equilibrium coincide. The key observation at this point is the following. Letting $T_i(x)$ denote i's optimal threshold when player j uses threshold x, we show that the mappings T_1, T_2 are contractions.

[11] There is equilibrium multiplicity in the sense that players are indifferent at their thresholds, but as long as $\theta_0 \notin \{\theta_*, \theta^*\}$, all equilibria yield the same outcome almost surely, since the probability that θ_t will ever equal θ_* or θ^* exactly is zero.

Thus there is a unique equilibrium in threshold strategies, which must be both the greatest and the smallest equilibrium.

Here is an intuition as to why T_1 is a contraction. Suppose that player 1's best response to a threshold θ_* is a threshold θ^* and her best response to a threshold $\theta_* + \varepsilon$ is a threshold $\theta' \geq \theta^* + \varepsilon$, where $\varepsilon > 0$. It can be shown that players must be indifferent at their thresholds, so equivalently, player 1 is indifferent about surrendering in state θ^* when facing threshold θ_*, and weakly prefers to continue in state $\theta^* + \varepsilon$ when facing threshold $\theta_* + \varepsilon$.

If we compare the resulting disputed regions in both scenarios, namely $[\theta_*, \theta^*]$ and $[\theta_* + \varepsilon, \theta']$, there is a clear contradiction. Indeed, in the latter scenario, player 1's payoff from continuing the war when at her threshold is worse for three reasons. First, since the new disputed region is made up of higher states, her flow costs over the course of the war are expected to be higher (Assumption B1). Second, because the stochastic process governing $(\theta_t)_t$ is not mean-reverting (Assumption A2), in the new disputed region, the drift of the state θ_t is weakly less favorable to her. Third, if $\theta' > \theta^* + \varepsilon$, the disputed region is larger, so it will take longer on average for the state to travel to player 2's surrender region.

From this discussion, it also follows that Assumptions A2 and B1 are jointly "tight" in the following sense: suppose that the $c_i(\theta)$ and F_θ are all constant in θ over some open interval $I \subseteq [-M, M]$, so that A2 is satisfied, and B1 is violated only slightly (i.e., there are cost functions arbitrarily close to the c_i which satisfy B1). Then, for any $\theta \in I$ such that $T_i(\theta) \in I$, $T_i'(\theta) = 1$. In addition, $T_1(\theta)$, $T_2(\theta) \in I$ for any $\theta \in I$ if we take H to be small enough. Hence, if H is small enough, then $(T_1 \circ T_2)'(\theta) = 1$ over some interval, which allows for multiple fixed points and hence multiple (indeed a continuum of) equilibria.[12] Thus we cannot relax B1 any further and still guarantee that Proposition 1 will hold, unless we tighten A2. Conversely, we can relax A2 only if we tighten B1.

We can interpret the thresholds θ_*, θ^* as parameterizing two features of the equilibrium. The size of the disputed region, $\theta^* - \theta_*$, reflects how willing the players are to fight to increase their odds of winning the war. The position of the disputed region, $[\theta_*, \theta^*]$ within the interval $[-M, M]$ reflects any asymmetries between the players. For example, if f and the flow costs are symmetric then $\theta^* + \theta_* = 0$, while if θ_t tends to drift to the right, or $c_1(\theta) > c_2(-\theta)$ (player 1 has higher costs), then $\theta^* + \theta_* < 0$, and so on.

Our next result characterizes the comparative statics of the model with respect to the prizes, the cost functions, and the stochastic process described by $(F_\theta)_{\theta \in [-M,M]}$. For cost functions, we say c_i is increased (decreased) if we replace c_i with a new cost function $\hat{c}_i$ such that $\hat{c}_i(\theta) > c_i(\theta)$ $(<)$ for all θ. For F_θ we apply the FOSD order: (weakly) increasing F_θ

[12] Alternatively, by analogous arguments, uniqueness is recovered if we require the c_i to be only weakly monotonic but require F_θ to be strictly FOSD-increasing.

means replacing F_θ with a $\hat{F}_\theta$ such that $\hat{F}_\theta(x) \leq F_\theta(x)$ for all x.

Proposition 2.

(i) Increases in H_1 and decreases in c_1 raise θ_*, θ^* and $\theta^* - \theta_*$. Increases in H_2 and decreases in c_2 lower θ_* and θ^* but raise $\theta^* - \theta_*$.

(ii) A weak increase in F_θ for all θ weakly lowers θ_* and θ^*.

Part (i) of Proposition 2 says that, if a player's prize from winning increases or her cost decreases, her surrender region shrinks, her opponent's surrender region expands, and the size of the disputed region expands. The logic is as follows: if H_1 goes up, for instance, this directly increases player 1's incentive to continue fighting, taking as fixed all the other parameters as well as player 2's strategy. Then player 1 shrinks her surrender region, which induces player 2 to expand her own. This incentivizes player 1 to shrink her surrender region further, and so on. Iterating this argument brings us to the new equilibrium thresholds. The new disputed region is made up of higher states—that is, states preferred by player 2. In order to leave player 2 indifferent at her new threshold—despite her now fighting over more favorable states—the size of the disputed region, $\theta^* - \theta_*$, must grow.

Part (ii) says that a similar logic applies to changes in the drift of the stochastic process $(\theta_t)_t$: making the evolution of the state more favorable to one player shrinks her surrender region and expands her opponent's.

The results in Proposition 2 can be translated into statements about changes in the players' winning probabilities and winning times. Formally, let $P_i(t)$ be the probability that player i will win by time t. Then Proposition 2 implies that, if H_i increases, then $P_i(t)$ increases for all t and $P_j(t)$ decreases for all t. Analogous statements hold for changes in c_i or F_θ.

Continuous Time Model

For much of the rest of the paper, we will study a special case of our model, in which taking the limit to continuous time is feasible and yields sharper results.

Assume that the underlying time index t is continuous, but that the players only make decisions at a discrete sequence of times: $t \in \{0, \Delta, 2\Delta, \ldots\}$. Let $\delta = e^{-\gamma}$ be the discount factor over a unit of time. Suppose that the state θ_t evolves according to a drift-diffusion process with reflecting boundaries at $-M$ and M, given by the expression

$$d\theta = \mu(\theta)dt + \sigma dB_t,$$

where $(B_t)_t$ is a Brownian motion, $\sigma > 0$ is fixed, and $\mu(\theta)$ is continuous (A1) and weakly increasing in θ (A2).[13] We can assume that costs $c_i(\theta_t)$ also accrue continuously, and let $\tilde{c}_i(\theta_t) = E\left[\int_t^{t+\Delta} e^{-\gamma(\tau-t)} c_i(\theta_\tau) d\tau | \theta_t\right]$ be the expected cost between t and $t+\Delta$ based on the interim evolution of the state.

Clearly, Propositions 1 and 2 apply for any value of $\Delta > 0$. As we take the limit $\Delta \to 0$, we converge to a continuous-time version of the model in which the players can surrender at any moment $t \in [0, +\infty)$; they discount the future at a common rate $\gamma \geq 0$; and, while the war continues, they pay flow costs $c_i(\theta_t)$ ($i = 1, 2$) satisfying Assumptions B1-5.[14]

Denote the best-response threshold functions from the previous section by T_i^Δ, highlighting the dependence on the time Δ between decisions. For any x, $T_i^\Delta(x)$ converges as $\Delta \to 0$ to a limit $T_i^0(x)$, i's optimal threshold in response to an opponent's threshold of x in the continuous-time game. And T_1^0, T_2^0 can be shown to be contractions using similar arguments. The limit equilibrium is then simply the fixed point of $T_1^0 \circ T_2^0$. For simplicity, we will denote the equilibrium disputed region in the continuous time setting simply as $[\theta_*, \theta^*]$. By our preceding arguments, $\theta_*^\Delta \to \theta_*$ and $\theta^{*\Delta} \to \theta^*$ as $\Delta \to 0$. Our next result provides an explicit characterization of the equilibrium thresholds, expected payoffs, and winning probabilities for both players in continuous time.

Proposition 3. *The unique equilibrium of the game with periods $t = 0, \Delta, 2\Delta, \ldots$ converges as $\Delta \to 0$ to an equilibrium of the continuous-time war of attrition. In it, the players use threshold strategies with thresholds $\theta_* < \theta^*$ such that each player's expected utility $V_i(\theta)$, conditional on an initial state θ, solves the following ODE:*

$$c_i(\theta) + \gamma V_i(\theta) = \mu(\theta) V_i'(\theta) + \frac{\sigma^2}{2} V_i''(\theta) \tag{1}$$

given the boundary conditions $V_1(\theta_*) = H_1$, $V_1(\theta^*) = 0$; $V_2(\theta_*) = 0$, $V_2(\theta^*) = H_2$; and the smooth-pasting conditions $V_1'(\theta^*) = 0$, $V_2'(\theta_*) = 0$.

[13] To guarantee A2 for low values of Δ, we require $\mu(\theta)$ to be weakly increasing and $\sigma(\theta)$ to be constant. This is due to the nature of normal noise. If $\sigma(\theta_1) > \sigma(\theta_0)$ for two states $\theta_1 > \theta_0$, then the distribution of marginal changes starting from θ_1 will not FOSD-dominate the other due to a fat left tail. Similarly, if $\sigma(\theta_1) < \sigma(\theta_0)$, it will not FOSD-dominate due to a thin right tail. However, these restrictions can be relaxed by a change-of-variables argument: if $\mu(\theta)$ and $\sigma(\theta)$ are such that $\frac{\mu(\theta)}{\sigma(\theta)} - \frac{\sigma'(\theta)}{2}$ is weakly increasing in θ, then, setting $h(\theta) = \int_0^\theta \frac{1}{\sigma(\tilde{\theta})} d\tilde{\theta}$, $\eta_t = h(\theta_t)$ has increasing drift and constant variance, as required.

[14] Assumption B4 must be restated as requiring that M be large enough that player 1's payoff from continuing the war starting at state M is negative, even when player 2 surrenders at all states $\theta \leq M_2$, and analogously for player 2. Also, Assumption B6 becomes vacuously true in continuous time.

Figure 1: Equilibrium utility and win prob.: $H_i = 2$, $c_1(\theta) = 5 + \theta$, $c_2(\theta) = 5 - \theta$, $\sigma^2 = 1$, $M = 15$

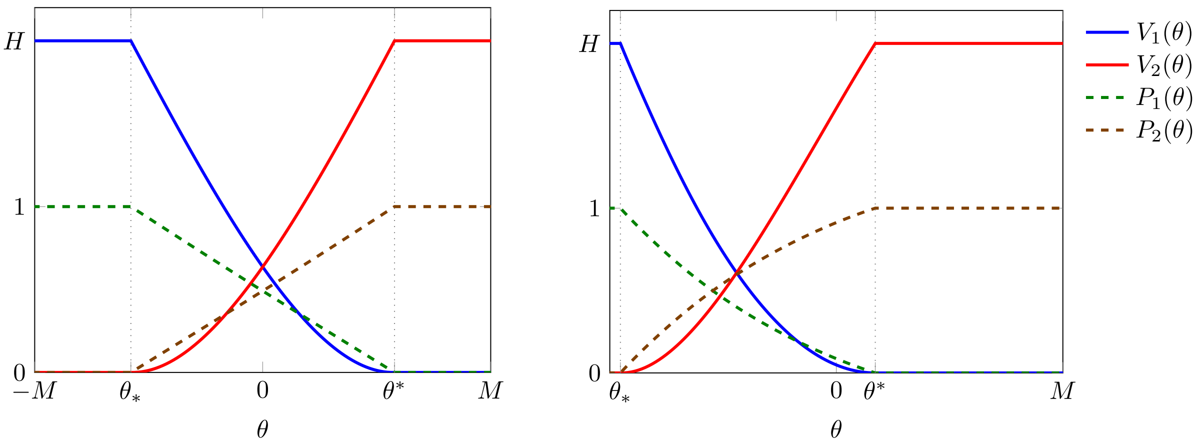

(a) $\mu(\theta) \equiv 0$

(b) $\mu(\theta) \equiv 0.0625$

The players' winning probability $P_i(\theta)$ solves the ODE:

$$0 = \mu(\theta)P_i'(\theta) + \frac{\sigma^2}{2}P_i''(\theta) \tag{2}$$

with boundary conditions $P_1(\theta_*) = 1$, $P_1(\theta^*) = 0$; $P_2(\theta_*) = 0$, $P_2(\theta^*) = 1$.

Equations 1 and 2 are obtained by combining the Hamilton-Jacobi-Bellman equation of each player's optimization problem with Itô's Lemma. In fact, *any* threshold strategy profile with thresholds $\underline{\theta} < \overline{\theta}$ induce value functions that solve Equation 1 with the boundary conditions $V_1(\underline{\theta}) = H_1$, $V_1(\overline{\theta}) = 0$; $V_2(\underline{\theta}) = 0$, $V_2(\overline{\theta}) = H_2$. It is the smooth-pasting conditions $V_1'(\theta^*) = 0$, $V_2'(\theta_*) = 0$ which serve to uniquely pin down the equilibrium thresholds θ_*, θ^*.

While Equations 1 and 2 do not have closed-form solutions in general, they simplify greatly if $\mu \equiv 0$ and $\gamma = 0$, i.e., when θ_t follows a Brownian motion with no drift and there is no discounting. In that case, Equations 1 and 2 reduce to:

$$V_1(\theta) = \frac{2}{\sigma^2}\int_\theta^{\theta^*}(\lambda - \theta)c_1(\lambda)d\lambda \qquad P_1(\theta) = \frac{\theta^* - \theta}{\theta^* - \theta_*} \tag{3}$$

$$V_2(\theta) = \frac{2}{\sigma^2}\int_{\theta_*}^{\theta}(\theta - \lambda)c_2(\lambda)d\lambda \qquad P_2(\theta) = \frac{\theta - \theta_*}{\theta^{*0} - \theta_*}, \tag{4}$$

plus the conditions $V_1(\theta_*) = H_1$, $V_2(\theta^*) = H_2$.

Figure 1 illustrates the expected utility and winning probabilities of the two players as a function of the initial value of θ when the stochastic process is symmetric ($\mu \equiv 0$, Figure 1a) and when it is asymmetric ($\mu > 0$, Figure 1b). In both examples, the cost functions are taken to be symmetric around 0. As expected, the thresholds and utilities in Figure 1a are symmetric around 0, and the winning probabilities $P_i(\theta)$ are linear in θ. On the other hand, in Figure 1b, θ tends to drift up over time, favoring player 2, and both players' thresholds are lower as a result.

The following Proposition shows that the comparative statics from Proposition 2 extend to the continuous time model (parts i and ii) and can be strengthened (part iii).

Proposition 4.

(i) Increases in H_1 and decreases in c_1 raise θ_*, θ^* and $\theta^* - \theta_*$. Increases in H_2 and decreases in c_2 lower θ_* and θ^* but raise $\theta^* - \theta_*$.

(ii) An increase in μ lowers θ_* and θ^*.

(iii) If H_1 and H_2 are increased proportionally, and $\mu(\theta) \geq 0$ for all θ, then θ_* decreases. Similarly, if $\mu(\theta) \leq 0$ for all θ, then θ^* increases. If $\mu(\theta) \equiv 0$, then θ_* decreases and θ^* increases.

Substantively, part (iii) of Proposition 4 states that, if both players' incentives to fight increase in the same proportion, the disputed region expands not just in the sense of $\theta^* - \theta_*$ growing—as stated in part (i)—but in the stronger sense that θ_* decreases while θ^* increases. The intuition is as follows. Suppose that, after increasing H_1 and H_2 proportionally, we are left in equilibrium with a new disputed region of the form $[\theta_*, \theta^* + \Delta]$ for some $\Delta > 0$, i.e., only player 1's threshold changes (the same argument applies if both thresholds move strictly in the same direction). Then all the states added to the disputed region are states in which player 2's costs are lower than player 1's, relative to the states in the original disputed region. Hence, if player 1 is indifferent at $\theta^* + \Delta$, player 2 must be strictly willing to continue the war at θ_*. The reason this result can only be proved in the continuous time setting is that the proof relies on the continuity of θ_t as a function of time—which guarantees that, in moving from one state to another, the path $(\theta_t)_t$ passes through every state between them.

The intuition in the previous paragraph applies when the process has no drift ($\mu \equiv 0$). If we instead have $\mu \geq 0$, i.e., the drift in θ_t favors player 2, a proportional increase in both players' prizes may either induce both thresholds to move away from each other, or it may induce both thresholds to move to the left, but they cannot both move to the right. The reason is that the player who faces a favorable drift has an intrinsic advantage that is amplified in longer wars. Thus, if prizes increase, the disputed region grows (part i), and this strengthens the position of player 2.

3 Discussion

Moving θ as Equilibrium Selection

The baseline model, with its time-varying state of the world θ_t, captures conflicts in which the players' strengths and weaknesses change meaningfully over time. However, the model can also be taken as a tool for equilibrium selection when the "true" model the researcher is interested in is the classic war of attrition. The correct object of study for this purpose is the limit of a sequence of equilibria as the movement of θ becomes arbitrarily slow.

In continuous time, this limit can be taken as follows. Denote $\tilde{\mu}(\theta) = \nu\mu(\theta)$, $\tilde{\sigma}(\theta) = \sqrt{\nu}\sigma$, and denote the equilibrium thresholds as a function of ν by $\theta_*(\nu)$, $\theta^*(\nu)$. We are interested in the limit $\nu \to 0$. Of course, the limit game obtained when $\nu = 0$ is the classic war of attrition, but taking the limit of the equilibrium disputed region $[\theta_*(\nu), \theta^*(\nu)]$ as $\nu \to 0$ uniquely selects an equilibrium of the classic war of attrition, described next.

Proposition 5. *Suppose $H_1 = H_2 = H$.[15] There is θ^l such that $\theta_*(\nu), \theta^*(\nu) \to \theta^l$ as $\nu \to 0$. If $\mu \equiv 0$, then θ^l is given by the condition $c_1(\theta^l) = c_2(\theta^l) =: c^*$, and as $\nu \to 0$, we have*

$$V_1\left(\frac{\theta_*(\nu) + \theta^*(\nu)}{2}\right), V_2\left(\frac{\theta_*(\nu) + \theta^*(\nu)}{2}\right) \longrightarrow \frac{H}{X + \frac{1}{X} + 2},$$

where $X = \sqrt{\frac{\gamma H}{c^*} + 1 + \sqrt{\left(\frac{\gamma H}{c^*} + 1\right)^2 - 1}}$.

Here is an explanation of this result. As $\nu \to 0$, the disputed region must shrink, as otherwise players would expect to wait a very long time before they can ever win the war, if the current state is near their surrender threshold. And, if there is no drift, the disputed region must contain the state of the world in which the players' cost functions intersect—else the disputed region would be composed entirely of states that favor player 1 over player 2 or vice versa, and it would be impossible for both players' indifference conditions to be met. Hence, as it shrinks, the disputed region converges to this point.

It follows that, if the initial state θ_0 is below θ^l, then the selected outcome in the classic war of attrition is that player 2 immediately surrenders. Similarly, if $\theta_0 > \theta^l$, player 1 immediately surrenders.

A different outcome arises when $\theta_0 \approx \theta^l$, i.e., when the players have equal prize-cost ratios. In this case, neither player surrenders immediately and both can win. Proposition 5 characterizes the welfare properties of this equilibrium in the limit: a significant fraction of total welfare is destroyed by fighting and delay, but not all. Indeed, when $\frac{H}{c}$ is small (i.e.,

[15] This normalization serves purely to simplify the notation.

the prize is small, so the players end the war quickly) or γ is small (players are patient, so delay does not in itself degrade the prize), $X \approx 1$, so $V_1, V_2 \approx \frac{H}{4}$, and $V_1 + V_2 \approx \frac{H}{2}$. In other words, total welfare is half of the first-best (where one player gets the prize immediately). In general, $X > 1$ and V_1, V_2 are below $\frac{H}{4}$, but still positive.

One implication of this result is that, as $\nu \to 0$, the equilibrium we select is *not* equivalent to the totally mixed symmetric equilibrium of the classic war of attrition, as that equilibrium leaves the players with zero rents. In fact, the equilibrium we select is not even a Nash equilibrium of the classic war of attrition; it is an equilibrium of the war of attrition *with a public coordination device*.[16]

We can think of the limit equilibrium as showing the players fighting over payoff-irrelevant *tokens*. In the war of attrition with tokens, players start with some number of tokens $K > 0$ and repeatedly play a payoff-irrelevant game that selects a random winner in each round; the winner takes a token from the loser in each round. The state θ_t tracks the amount of tokens held by player 2. The first player to run out of tokens surrenders.

Alternative Models

This Section discusses the predictions of the model in comparison to those made by other variants of the war of attrition found in the literature. We begin with the classic war of attrition.

As discussed above, the classic war of attrition is nested in our model when $\mu \equiv \sigma \equiv 0$, i.e., θ is constant. (Or, we can take $c_1(\theta) \equiv c_1$ and $c_2(\theta) \equiv c_2$ to be flat.) For simplicity we normalize $H_1 = H_2 = H$.

It is known (Hendricks et al., 1988) that this game has a continuum of subgame perfect equilibria, which are as follows. In every equilibrium, at all times $t > 0$, each player i surrenders at a rate $\frac{c_j}{H}$. The equilibria differ in what happens at $t = 0$. For each $p \in [0, 1]$ and each player i, there is an equilibrium in which i surrenders with probability p at $t = 0$, while j does not surrender at $t = 0$.

In particular, there is an equilibrium in which player 1 surrenders immediately; another in which player 2 surrenders immediately; and a mixed strategy equilibrium with no instantaneous concession. In the latter, both players' expected payoffs are 0, and i's probability of winning is $\frac{c_i}{c_i + c_j}$.[17] In every equilibrium, at least one player has an expected payoff of 0.

The results in this paper differ from the above in several ways. The model in this paper

[16] But also not the most efficient equilibrium with a public coordination device. Indeed, the players could have a coin toss and instantly have the loser concede, attaining a fair outcome with no welfare loss.

[17] Strangely, in this equilibrium, a player's probability of winning increases with her *own* cost. Holding c_j constant, in the limit as $c_i \to \infty$, player i wins almost surely.

has a unique equilibrium. Assuming an initial state in the disputed region, this equilibrium gives both players a positive expected payoff, even when payoff perturbations are small (Proposition 5), which is impossible in the classic model. As for expected delay, the classic war of attrition admits equilibria with any expected delay between zero and the maximum (i.e., enough to completely evaporate the value of the prize). The model in this paper predicts an amount of delay strictly between these extremes, as implied by Proposition 5.

The two models also predict delay distributions with different shapes. In the classic model, the surrender rate is constant, and there may be immediate surrender. Hence the density $h(t)$ of the length of the war is exponentially decreasing over time, with a possible spike at $t = 0$ (Figure 2a). In contrast, in the model presented in this paper, there is no immediate surrender. Moreover, in the discrete time version, the length of the war is bounded away from zero: any time T at which the war ends must satisfy $T \geq \frac{\min(\theta^* - \theta_0, \theta_0 - \theta_*)}{\eta} > 0$. In continuous time, $\theta_t - \theta_0$ has full support for all $t > 0$, owing to properties of the Brownian motion. But it is still approximately true: as shown in Figure 2b, $h(t) \to 0$ as $t \to 0$, as it typically takes some time for θ to hit either surrender threshold. Thus, the distribution of delay is hump-shaped over time. This difference may be used to distinguish between the two models empirically.

Figure 2: Density of length of the war

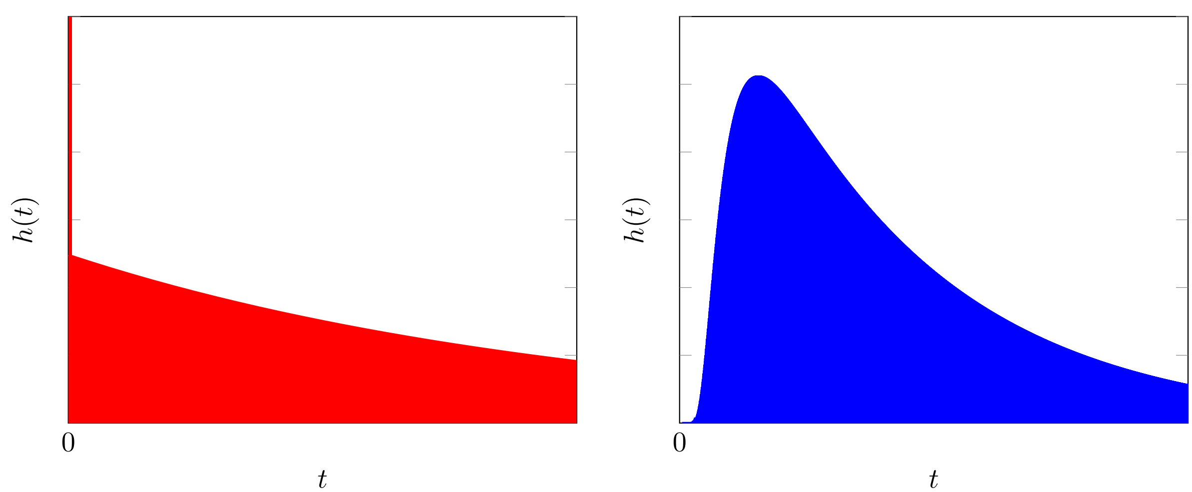

(a) Classic/reputational war of attrition (b) Evolving war of attrition

Next, consider the war of attrition with reputational concerns. The simplest version (Abreu and Gul, 2000) is identical to the classic war of attrition, except that each player i has an exogenous probability ϵ_i of being a *commitment* type that never surrenders. The game has a unique equilibrium, which is observationally equivalent to one of the equilibria of the classic war of attrition described above (up to the point where only comitment types are left). The innovation is that the probability of immediate surrender (and who surrenders) are uniquely determined as a function of the ϵ_i. Much of the comparison made above with the classic war of attrition thus extends to the Abreu and Gul (2000) setting: indeed, in Abreu and Gul (2000), (the rational types of) at least one player must have an expected payoff of 0, and the distribution of delay must have exponential density with a spike at $t = 0$, unlike in this paper.

In other variants of the reputational war of attrition (e.g., Fudenberg and Tirole (1986)), each player has a continuum of possible flow costs. If each player has a positive probability of having negative cost—meaning she will never surrender—the game has a unique equilibrium, featuring smooth screening of the types with positive costs, possibly preceded by the immediate surrender of some types of one player. In such a model, (most) types of each player have positive expected payoffs, but the distribution of delay differs from our model's. Indeed, in Fudenberg and Tirole (1986), both players have a positive surrender rate at all times $t > 0$, much like in Abreu and Gul (2000). Moreover, if we "take the perturbation to zero" (i.e., we collapse the distribution of each player's flow costs towards a mass at c), the ex ante payoff of at least one player goes to zero, again as in Abreu and Gul (2000), because the equilibrium must converge to an equilibrium of the classic war of attrition. What sets our model apart is thus that we do **not** select an equilibrium of the classic war of attrition, even in the limit.

Finally, we can compare our setting to Gul and Pesendorfer (2012)'s war of information. Their baseline model is equivalent to a special (limiting)[18] case of our continuous time model with the following parameters: $c_1(\theta) = k_1 \mathbb{1}_{\theta>0}$, $c_2(\theta) = k_2 \mathbb{1}_{\theta<0}$, $\gamma = 0$, $\sigma(\theta) \equiv 1$ and $\mu(\theta) = -\frac{1}{2} + \frac{1}{1+e^{-\theta}}$.[19] They then extend their results to allow for $\gamma > 0$, and separately for a specific form of variable costs: $c_1(\theta) = k_1 \theta \mathbb{1}_{\theta>0}$, $c_2(\theta) = -k_2 \theta \mathbb{1}_{\theta<0}$ (with $\gamma = 0$). They provide a result analogous to Proposition 1. While their arguments regarding supermodularity and extremal equilibria in threshold strategies are analogous to the ones given here, their proof that $T_1 \circ T_2$ is a contraction, as required to show uniqueness, relies on explicitly calculating the players' utilities from an arbitrary strategy profile, and from there deriving explicit expressions for $(T_1 \circ T_2)'$ which are shown by hand to be less than 1 in each case they study. This approach does not generalize well to general cost and drift functions. Our argument shows that $(T_1 \circ T_2)'$ is a contraction under much less restrictive conditions, and for transparent reasons. They also consider a limit of their game as news become uninformative (see their Proposition 2); their model, however, does not converge to the classic war of

[18] The cost functions violate Assumption B1, but our results go through as μ_θ is strictly increasing; see footnote 12.

[19] This equivalence is obtained by setting $\theta_t = \ln\left(\frac{p_t}{1-p_t}\right)$, where p_t is the probability that party 1 is better given the information revealed at time t, as defined in their paper.

attrition in this limit, as they retain the assumption of discontinuous cost functions (in particular, only one player pays a flow cost at each state). This distinguishes their results from our Proposition 5.

4 Partial Concessions

In the baseline model, as in the classic war of attrition, the players only have two choices at each moment: continue or surrender completely. In other words, they can control the duration of the war, but not its intensity. To illustrate what is being ruled out, take the example of a polluting firm being boycotted by an activist group. In the baseline model, the war ends when the firm capitulates to all the demands or the activists abandon the boycott; there is no room for an intermediate solution. However, in practice the firm may have access to a range of policies it can implement to lower its own pollution. It may prefer to announce a *partial concession*, in the form of a unilateral commitment to a certain level of self-regulation. Such an announcement may deflate the boycott's momentum even if it does not fully meet the activists' demands. A similar logic applies, e.g., to a government facing protests.

This Section introduces an extension of the model that allows for partial concessions. In other words, it allows for surrender at both the intensive and extensive margins. To illustrate the relevant forces as simply as possible, we will focus on the case where only one player has the ability to make concessions. Such a restriction is reasonable when modeling collective action: for a boycott or protest movement, announcing that they will give up on one of their demands may be mechanically or reputationally difficult. It is also appropriate, for instance, in a siege: while a besieged city may offer some tribute to an attacking army to encourage it to leave, the army cannot credibly commit to leaving after the tribute has been received, or even to not fully destroying the city if the gates are later opened. (We consider two-sided concessions in Appendix C.)

The general insight that emerges is that unilateral concessions can resolve conflict, and can also be to the conceder's advantage, but are useful *only* when the prize is composed of heterogeneous parts that are unequally valued by the two players. To accommodate this possibility, we represent the prize as an interval $[0, 1]$, where $v_i(x)$ is the value that i assigns to part x of the prize. We assume that v_1 is weakly decreasing in x, v_2 is weakly increasing in x, and denote $H_1 = \int_0^1 v_1(x)dx$, $H_2 = \int_0^1 v_2(x)dx$.[20]

[20] A simpler specification might be that, if the prize is split in shares $(x, 1-x)$, then payoffs from the prize are xH_1, $(1-x)H_2$. This is the case of a homogeneous prize, in which, as we will see, concessions are not useful.

Formally, we begin by considering the following game. At $t = 0$, player 1 can choose a cutoff x^* and unilaterally give up everything to the right of x^*, so that player 2 collects a payoff $\int_{x^*}^1 v_2(x)dx$ immediately.[21] Afterwards, further concessions are not possible, and a war of attrition is played over the remaining prize $[0, x^*)$.

We assume an especially tractable specification of the war of attrition: we assume as in Proposition 3 that, after the concession stage, time is continuous,[22] and in addition $\mu \equiv 0$ and $\gamma = 0$, i.e., θ_t follows a Brownian motion with no drift and there is no discounting. The continuation game is thus identical to that from Proposition 3, but with prizes $\tilde{H}_1 = \int_0^{x^*} v_1(x)dx$, $\tilde{H}_2 = \int_0^{x^*} v_2(x)dx$.

We denote the equilibrium disputed region in this continuation by $[\theta_*(x^*), \theta^*(x^*)]$. From applying Equations 3 and 4, we know that these thresholds uniquely solve the following system:

$$\tilde{H}_1 = \int_0^{x^*} v_1(x)dx = V_1(\theta_*(x^*)) = \frac{2}{\sigma^2} \int_{\theta_*(x^*)}^{\theta^*(x^*)} (\lambda - \theta_*(x^*))c_1(\lambda)d\lambda \tag{5}$$

$$\tilde{H}_2 = \int_0^{x^*} v_2(x)dx = V_2(\theta^*(x^*)) = \frac{2}{\sigma^2} \int_{\theta_*(x^*)}^{\theta^*(x^*)} (\theta^*(x^*) - \lambda)c_2(\lambda)d\lambda. \tag{6}$$

The model we have set up makes two strong assumptions. First, concessions can only be made at the beginning of the game. Second, concessions are unilateral, meaning that the receiver *cannot* refuse them, even if the resulting equilibrium gives her a lower payoff than she would get in the baseline model with no concessions. We discuss the possibility of concessions made at later times, as well as concessions that the receiver can veto, after the main results.

The following Proposition provides a characterization of equilibrium.

Proposition 6. *An equilibrium exists. In any equilibrium, player 1 chooses a value of x^* that solves:*

$$max_{x \in [0,1]} \theta^*(x) \quad \textit{subject to } \theta_*(x) \leq \theta_0.$$

If this maximization problem has a unique solution, then the equilibrium is essentially unique.

In other words, player 1 makes the concession that makes her most aggressive in the continuation, by maximizing her own surrender threshold $\theta^*(x)$—with the caveat that it is never useful to *strictly* induce player 2 to surrender immediately. To see why, note that, if player 1's concession does *not* induce immediate surrender by her opponent (that is, $\theta_*(x^*) \leq \theta_0 \leq \theta^*(x^*)$), then her value function over $[\theta_0, \theta^*(x^*)]$ can be calculated by solving Equation

[21] It would be worse for player 1 to concede a set not of the form $[x^*, 1]$, if this were allowed.

[22] As in Proposition 3, the equilibria we characterize are limits of discrete time equilibria for $\Delta \to 0$.

1 leftwards from $\theta^*(x^*)$, with the initial conditions given by smooth-pasting. As the solution is increasing in the starting point $\theta^*(x^*)$, player 1 need only make $\theta^*(x^*)$ as high as possible. An implication that is perhaps surprising is that, for any θ_0 between $\theta_*(x^*)$ and $\theta^*(x^*)$, the optimal concession is independent of the initial state θ_0.

We can provide a geometric description of the optimal concession x^*. For each possible value of $\theta^* \in [-M, M]$, consider all pairs (H_1, H_2) that make 1's surrender threshold equal to θ^* per Proposition 3. In this way, partition the space of possible prize pairs (H_1, H_2) into level curves, parameterized thus: $H_2(H_1; \theta^*)$. The concession technology yields a feasible path of prize pairs $(H_1(x), H_2(x)) = (\int_0^x v_1(\tilde{x})d\tilde{x}, \int_0^x v_2(\tilde{x})d\tilde{x})$. Player 1 then picks the prize pair on the highest possible level curve, as illustrated in Figure 3.

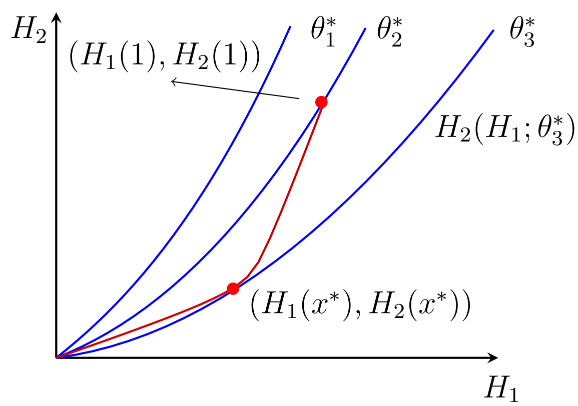

Figure 3: Level curves, feasible prize pairs, and optimal concession x^*

Remark 1. H_2 is a convex function of H_1 along all level curves, as well as along the feasible path of prize pairs.

The convexity of the prize path is due to the shape of the functions v_i, while the convexity of the level curves follows from Proposition 4.(iii). If x^* is interior (i.e., $x^* < 1$ and it does not induce immediate surrender), then the path of prize pairs must be tangent to the level curve passing through x^*, so the slopes must be equal. From Equations 3–4, and since the slope of the path of prize pairs is $\frac{v_2(x)}{v_1(x)}$, we have

$$\frac{c_2(\theta_*(x^*))}{⨍_{\theta_*(x^*)}^{\theta^*(x^*)} c_1(\lambda)d\lambda} = \frac{v_2(x^*)}{v_1(x^*)}. \tag{7}$$

A derivation is given in the proof of Remark 1.

Our analysis yields sharp predictions in two special cases. First, suppose that the prize is homogeneously valued by both players, that is, $v_1(x)$ and $v_2(x)$ are constant in x. Then the path of feasible prize pairs is simply a line segment from $(0, 0)$ to (H_1, H_2), and the highest level curve is attained at $x^* = 1$, i.e., by making no concession. The same result holds if the v_i are approximately constant. In other words, player 1 can only benefit from making concessions when the valuations of the conceded prize are lopsided enough that the gain from undermining player 2's incentives to fight over the remainder dominates the direct loss from foregoing part of the prize. Second, suppose that $v_1(x) \equiv \overline{v}_1 > 0$ is constant, while $v_2(x) \equiv 0$ for $x \leq \overline{x}$ and $v_2(x) \equiv \overline{v}_2 > 0$ for $x > \overline{x}$. In other words, part of the prize is homogeneous, and part is valued only by player 1.[23] In this case the relevant part of the feasible path of prize pairs is a line segment from $(\overline{x}\overline{v}_1, 0)$ to (H_1, H_2). Since an interior optimum is impossible, the optimal concession is either no concession at all or the minimal concession that induces immediate surrender. However, in the general case, the optimal concession may be positive while not immediately ending the war.

It is worth comparing these results to those of a standard bargaining setting. For instance, if player 1 gets to make a single take-it-or-leave-it offer, she would offer $[\hat{x}, 1]$, where $\hat{x}$ is chosen to leave player 2 indifferent about accepting, i.e., $\int_{\hat{x}}^1 v_2(x)dx = V_2(\theta_0)$, and the war would end immediately. (If there are alternating offers, a lower $\hat{x}$ would be chosen that splits the bargaining surplus more equally.) In contrast, the equilibrium concession x^* in our model does not necessarily end the war; it may be higher or lower than $\hat{x}$; and it may even leave player 2 worse off than she is either in the bargaining setting or when concessions are impossible. The reason is that unilateral concessions differ from bargaining offers in two ways. On the one hand, player 2 can always continue to fight after a concession, i.e., there is no chance of a *quid pro quo*, so player 1 may need to make a larger concession to deplete player 2's incentives to fight—or, if the effect on player 2's incentives is not strong enough, she may concede nothing at all. Second, player 2 cannot veto a "stingy" concession ($x^* > \hat{x}$); rather than refusing the offer, her only recourse is to continue fighting. To see how this may benefit player 1, suppose for instance that θ_t is slow-moving, as in Proposition 5, and θ_0 is only slightly higher than θ^l. Then player 1 would lose immediately in the absence of a concession, and so would have to give up the entire prize in bargaining; but, if the v_i are not flat, she can make herself the stronger player with a very small concession, after which player 2 will be the one forced to surrender.

The assumption that concessions are unilateral is reasonable in the applications that we discussed. For instance, an activist cannot *prevent* a firm from adopting a half-hearted level of self-regulation in response to a boycott. They can announce that such a concession is not enough, that it is meaningless, etc., but this is cheap talk; the effect on payoffs remains. However, in other settings, player 2 may get to veto an offered concession that is

[23]For instance, if player 1 is a city under siege, both players may value the city's wealth, but only the city values the lives of its citizens.

strategically damaging. For example, in January 2024, Senate Republicans backed out of a bill restricting immigration that was offered as a concession by Democrats in exchange for Ukraine funding, because Trump wanted to keep immigration active as a campaign issue for the 2024 election.[24] (While a presidential election is not a war of attrition, it is a type of race (Harris and Vickers, 1987), and thus closely related.)

The possibility of a veto by player 2 restricts player 1's set of feasible concessions to all $x \in [0, 1]$ such that $V_2(\theta_0; \tilde{H}_1(x), \tilde{H}_2(x)) + \int_x^1 v_2(y)dy \geq V_2(\theta_0; H_1, H_2)$. From this restricted (closed) set, player 1 will still choose either the lowest element inducing immediate surrender or the one maximizing $\theta^*(x)$. If the unrestricted optimum from Proposition 6 is incentive-compatible for player 2, then the solution is unchanged and the veto has no bite. When the veto has bite, player 1's equilibrium concession may grow or shrink relative to Proposition 6. Moreover, using Equation 1, it can be shown that $V_2(\theta_0; H_1, H_2) - V_2(\theta_0; \tilde{H}_1(x), \tilde{H}_2(x))$ is increasing in θ_0, i.e., the stronger player 2's initial position, the higher the strategic cost of allowing a concession. In particular, when player 2 would almost certainly lose absent concessions, almost any concession is acceptable, while when he would likely win absent concessions, any attempted concession gets rejected. Player 1's offer must then condition more carefully on θ_0.

Finally, we consider the case when player 1 can make concessions at any time, including multiple concessions over time. To avert some technical issues, we assume that the prize can only be divided into a finite number of pieces, i.e., there is a sequence $0 = x_0 < x_1 < \ldots < x_k = 1$ such that only concession thresholds $x = x_i$ are feasible, with $x^* = x_i$ for some i; and player 1 can concede from fighting over $[0, x_j)$ to fighting over $[0, x_l)$ for any $l < j$ at any time t, while player 2 can only continue or surrender. Then we have the following result:

Proposition 7. *Let $x^* := \arg\max_{x \in [0,1]} \theta^*(x)$ be the global maximizer of $\theta^*(x)$, and suppose it is unique. Furthermore, assume $\theta_*(x^*) < \theta_0$. Then the equilibrium outcome in Proposition 6 is also the equilibrium path for an equilibrium of this game. Moreover, player 1 cannot obtain a higher payoff in any threshold strategy equilibrium.*

Intuitively, the timing of concessions makes no difference because there is no discounting, and the optimal concession is largely independent of the current state, so there is little temptation to adjust it as the war progresses. However, equilibria in which player 1's equilibrium payoff is strictly lower than when concessions are restricted to $t = 0$ *can* exist. The logic is as follows: when all concessions by player 1 weaken her position, she chooses to concede nothing in the one-shot case—and, by doing so, she can *commit* not to conceding

[24] `https://www.cnn.com/2024/01/25/politics/gop-senators-angry-trump-immigration-deal/index.html`

in the future, unless she surrenders. However, when the option to make a concession in the future is always open, player 2 may always expect such a concession, and hence may have a higher (expected) continuation value and a lower surrender threshold before a concession, whence player 1 may choose to concede. Such equilibria disappear if we assume that partial concessions can only be made for a limited time (i.e., for $t \leq t_0$), or if player 1 can commit to not making partial concessions in the future.

5 Multi-Dimensional State

The baseline model imposes that the state must affect the players' payoffs in opposite ways. That is, θ is one-dimensional, and a higher θ is good for player 2 *and* bad for player 1. This assumption simplifies the analysis but is not essential. Indeed, this Section extends the model to allow for a multi-dimensional state θ_t, in which an analog of Proposition 1 holds. The substantive upshot is that the multi-dimensional model can capture richer variations in the players' costs. For example, in a price war between two firms, θ_t can be two-dimensional, with one dimension representing total demand and the other representing relative market share. Additionally, the stochastic process governing the state need only be "mean-reinforcing" in one direction, so the multi-dimensional model can accommodate shocks that are mean-reverting, cyclical, etc. by incorporating them into the additional dimensions of the state.

As in Section 2, there are two players living in discrete time with infinite horizon. In each period, each player can choose to continue or surrender. There is a state of the world $\theta_t \in \mathcal{M} = \mathbb{R} \times \prod_{i=1}^{k}[-M_i, M_i]$ which is common knowledge at all times. The initial θ_0 is a parameter. For $t > 0$, it evolves according to a Markov process:

$$P(\theta_{t+1} - \theta_t \leq x | \theta_t) = F_{\theta_t}(x),$$

where, for each θ, $F_\theta : \mathcal{M} \to [0, 1]$ is an absolutely continuous joint c.d.f. with density f_θ, and $x \leq y$ iff $x_i \leq y_i$ for all i.

Before proceeding we will need two definitions. Denote $v = (1, 0, \dots, 0) \in \mathbb{R}^{k+1}$. First, given two distributions G, H over $\mathcal{M}$, we say that *G FOSDs H* if there is a probability space $(\Omega, \mathcal{F}, P)$ in which there exist random variables X, Y, $Z : \Omega \to \mathbb{R}^{k+1}$ such that $X \sim G$, $Y \sim H$, $X = Y + Z$, and $Z(\omega) \equiv \alpha(\omega)v$ for a non-negative random variable $\alpha : \Omega \to \mathbb{R}_{\geq 0}$. That is, X always differs from Y by a weakly positive multiple of v. Second, we will say that a set $A \subseteq \mathcal{M}$ is *monotonic* if $\theta \in \mathcal{M} \implies \theta + av \in \mathcal{M}$ for all $a > 0$. Analogously, A is *antimonotonic* if $\theta \in \mathcal{M} \implies \theta + av \in \mathcal{M}$ for all $a < 0$.

We will assume that, for some $\eta > 0$,

A1' f_θ is continuous in θ, i.e., the mapping $\theta \mapsto f_\theta$ is continuous, taking the 1-norm in the codomain.

A2' F_θ is weakly FOSD-monotonic in θ for all $\theta \in \mathcal{M}$, in the following sense: if θ, θ' are such that $\theta = \theta' + av$ for some $a > 0$, then F_θ FOSDs $F_{\theta'}$.

A3' For all $\theta \in \mathcal{M}$, $\text{supp}(F_\theta)$ is a convex compact set with nonempty interior such that $0 \in \text{supp}(F_\theta) \subseteq B(0, \eta) \cap \left(\mathbb{R} \times \prod_{i=1}^{k}[-M_i - \theta_i, M_i - \theta_i]\right)$.

Assumptions A1' and A3' are natural adaptations of the analogous assumptions given in Section 2. The appropriate version of A2—ruling out mean reversion—is less obvious. Intuitively, A2' says that, if a state θ is higher than another state θ' (in the sense of being higher in the first dimension), then the drift of the state conditional on starting at θ is also a shifted-up version of the drift starting at θ'. (In particular, A2' is automatically true if f_θ is equal to a fixed density f for all θ, or more generally if the first dimension of θ_t satisfies A2 and evolves independently of the other dimensions.)

We assume the following about the players' payoff functions:

B1' $c_1(\theta)$ is strictly increasing, and $c_2(\theta)$ is strictly decreasing, in the first argument of θ. That is, if $\theta = \theta' + av$ for $a > 0$, then $c_1(\theta) > c_1(\theta')$ and $c_2(\theta) < c_2(\theta')$.

B3' There is a finite $D > 0$ such that, for any $\theta \in \mathcal{M}$, there are $a < b \in \mathbb{R}$ such that $c_1(\theta + xv) \leq 0$ for all $x \leq a$, $c_2(\theta + xv) \leq 0$ for all $x \geq b$, and $b - a \leq D$.

plus B2, B5 and B6, which are unchanged from the baseline model.

Again, B1' and B3' are analogous to B1 and B3, respectively. B1' now requires that shifting the state up in the first dimension is bad for player 1 and good for player 2, while B3' requires that high enough positive (negative) shifts in the first dimension take player 2 (1) into a region where fighting yields a benefit rather than a cost.[25] No analog of B4 is needed as we have effectively taken $M = \infty$ in the relevant (first) dimension.

The following Proposition characterizes the equilibrium of this game.

Proposition 8. *There is a unique SPE. In it, player 1 surrenders whenever $\theta_t \in A$, and player 2 surrenders whenever $\theta_t \in B$, where A, $B \subset \mathcal{M}$ are disjoint sets such that A is monotonic and B is antimonotonic.*

This generalized model nests the following natural example. Consider a price war between two duopolists. Let $\theta_t = (\theta_{1t}, \theta_{2t})$. Assume that $c_1(\theta)$ is strictly increasing in θ_1 and $c_2(\theta)$

[25] For technical reasons, B3' also requires that the gap between the players' dominance regions be uniformly bounded. Our results hold for any D, so this condition can be relaxed as much as desired.

is strictly decreasing in θ_1, while both flow costs are decreasing in θ_2. We can interpret θ_1 as representing player 2's market share, and θ_2 as representing total demand. (Note that A2' is satisfied automatically if we assume that $f_\theta \equiv f$ is independent of θ.) Thus, both the distribution of the pie as well as the size of the pie, whereas the baseline model essentially only allowed the players to steal market share from each other, with the market size held fixed.

The main properties of the equilibrium found in Proposition 1 are still present here. The equilibrium is unique, and is given by three sets: a surrender region for each player, and the disputed region separating them. If the initial state is in the disputed region, both players have a positive probability of winning, and they both have positive expected payoffs. In addition, Proposition 2 can be extended to this model: increasing H_1 will shrink player 1's surrender region and expand player 2's, etc.

Among other things, this extension is useful in comparing this paper to Georgiadis et al. (2022). Georgiadis et al. (2022) studies a war of attrition with a time-varying, one-dimensional state θ_t which affects both players symmetrically, rather than in opposite ways. Player 1 is assumed to have a lower outside option; the players are otherwise identical.[26] The authors show that, in every (Markovian) equilibrium of their model, there is a player who is guaranteed to quit first, regardless of how the state evolves; but who this player is may depend on the equilibrium. More precisely, there is always an equilibrium in which player 1 never surrenders before player 2. If the outside options differ by enough, this is the only equilibrium; else there are also equilibria in which player 2 never surrenders before player 1.

These results contrast with the ones in this paper. Indeed, the equilibrium may not be unique; in every equilibrium, there is no uncertainty about who quits first; and, consequently, the player who surrenders first plays as if she expected the opponent to never surrender. In contrast, in this paper, both players "gamble" on the other player surrendering first.

The model in this Section—in particular, the duopoly example—approximately nests both Georgiadis et al. (2022) and the main model from Section 2. Setting θ_{1t} constant yields the model in Georgiadis et al. (2022), while setting θ_{2t} constant yields our baseline model. The multi-dimensional model allows us to consider any (imperfect) correlation structure, and hence can approximate both extremes.

More precisely, let $\theta_{1t} = \mu dt + \sigma dB_t$ and $\theta_{2t} = \tilde{\mu} dt + \tilde{\sigma}\tilde{B}_t$, where B_t, $\tilde{B}_t$ are independent Brownian motions. Then setting $\tilde{\mu} = \tilde{\sigma} = 0$ yields a degenerate case analogous to our baseline model, whereas setting $\mu = \sigma = 0$ yields a degenerate case analogous to Georgiadis et al. (2022), in which their main results hold. Proposition 8 shows that the main results from our baseline model survive for $\sigma, \tilde{\sigma} > 0$. Moreover, it can be shown that, as $\sigma, \mu \to 0$,

[26] This is equivalent to player 1 having a higher prize and lower costs.

the selected equilibrium converges to one in which player 2—the one with the higher outside option—always exits first. Thus, adding slight uncertainty about the players' future relative strengths allows us to eliminate the equilibrium multiplicity in Georgiadis et al. (2022).

6 Conclusions

We have shown that the addition of an evolving state of the world to the classic war of attrition yields several attractive properties absent from the unperturbed game: the equilibrium is unique and the comparative statics are well behaved. In particular, if a player's prize increases or her cost decreases, she is more likely to win, and the war will end sooner if she held an advantage to begin with, whereas it will lengthen if she was an underdog at first. The model can be augmented to allow for partial concessions. The logic that arises is that concessions can benefit the conceder when they disproportionately sap the opponent's incentive to fight.

Relative to models with reputational concerns, the evolving war of attrition makes predictions that are less sensitive to small perturbations to the parameters (in particular, the players' reputations). The two frameworks diverge further when concessions are allowed: since our model has complete information, the players need not worry that concessions will signal weakness, as they might in a reputational model.

The logic of unilateral concessions is also distinct from that of offers in a bargaining framework: concessions cannot be rejected, but they also cannot be made in exchange for a matching concession from the opponent. As a result, it is harder to end a war when only unilateral concessions are available, and the conceder may do better or worse than in a bargaining setting.

Other applications are possible. For instance, the game can be extended to include costly commitment devices, i.e., "bridge-burning"; 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 (e.g., firms may engage in a price war with limited scope, or countries may ban certain types of weapons to limit the costs of war). Such actions are commonplace, and have been discussed by game theorists since at least Schelling (1960), but not given a systematic treatment within the war of attrition framework.

References

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

Alesina, Alberto and Allan Drazen, "Why Are Stabilizations Delayed?," *American Economic Review*, 1991, *81* (5), 1170–1188.

Bliss, Christopher and Barry Nalebuff, "Dragon-Slaying and Ballroom Dancing: The Private Supply of a Public Good," *Journal of Public Economics*, 1984, *25* (1-2), 1–12.

Budd, Christopher, Christopher Harris, and John Vickers, "A model of the evolution of duopoly: Does the asymmetry between firms tend to increase or decrease?," *The Review of Economic Studies*, 1993, *60* (3), 543–573.

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.

Cao, Dan, "Racing under uncertainty: Boundary value problem approach," *Journal of Economic Theory*, 2014, *151*, 508–527.

Egorov, Georgy and Bård Harstad, "Private Politics and Public Regulation," *The Review of Economic Studies*, 2017, *84* (4), 1652–1682.

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

_ and _ , *Game Theory*, Cambridge, MA: MIT Press, 1991.

Georgiadis, George, Youngsoo Kim, and H Dharma Kwon, "The Absence of Attrition in a War of Attrition under Complete Information," *Games and Economic Behavior*, 2022, *131*, 171–185.

Gul, Faruk and Wolfgang Pesendorfer, "The War of Information," *The Review of Economic Studies*, 2012, *79* (2), 707–734.

Harris, Christopher and John Vickers, "Perfect equilibrium in a model of a race," *The Review of Economic Studies*, 1985, *52* (2), 193–209.

_ and _ , "Racing with Uncertainty," *The Review of Economic Studies*, 1987, *54* (1), 1–21.

Hendricks, Ken, Andrew Weiss, and Charles Wilson, "The War of Attrition in Continuous Time with Complete Information," *International Economic Review*, 1988, pp. 663–680.

Kapur, Sandeep, "Markov Perfect Equilibria in an N-Player War of Attrition," *Economics Letters*, 1995, *47* (2), 149–154.

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

Maskin, Eric and Jean Tirole, "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cdycles," *Econometrica: Journal of the Econometric Society*, 1988, pp. 571–599.

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.

Montez, João, "Inefficient Sales Delays by a Durable-Good Monopoly Facing a Finite Number of Buyers," *The RAND Journal of Economics*, 2013, *44* (3), 425–437.

Morris, Stephen and Hyun Song Shin, "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks," *American Economic Review*, 1998, pp. 587–597.

Moscarini, Giuseppe and Lones Smith, "Optimal Dynamic Contests," 2011.

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.

Pitchford, Rohan and Mark LJ Wright, "Holdouts in Sovereign Debt Restructuring: A Theory of Negotiation in a Weak Contractual Environment," *The Review of Economic Studies*, 2012, *79* (2), 812–837.

Powell, Robert, "Bargaining and Learning while Fighting," *American Journal of Political Science*, 2004, *48* (2), 344–361.

Schelling, Thomas C, *The Strategy of Conflict*, Harvard University Press, 1960.

Slantchev, Branislav L, "The Power to Hurt: Costly Conflict with Completely Informed States," *American Political Science Review*, 2003, *97* (1), 123–133.

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

Smith, Alastair, "Fighting Battles, Winning Wars," *Journal of Conflict Resolution*, 1998, *42* (3), 301–320.

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.

Wang, Zhongmin, "(Mixed) Strategy in Oligopoly Pricing: Evidence from Gasoline Price Cycles Before and Under a Timing Regulation," *Journal of Political Economy*, 2009, *117* (6), 987–1030.

Winters, Harold A, *Battling the Elements: Weather and Terrain in the Conduct of War*, JHU Press, 2001.

A Proofs

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

We will say histories to refer to non-terminal histories for brevity, and write $h_t = (\theta)^t$ as shorthand for $h_t = ((\theta)^t, (1, \ldots, 1), (1, \ldots, 1))$. We will denote by H the set of all (non-terminal) histories and by H_t the set of all (non-terminal) histories at time t.

Definition 2. A *strategy* for player i, denoted by ψ_i, is a collection of probabilities $\psi_i(h) \in [0, 1]$ for each history h, denoting the probability that player i continues at history h.

We assume that players can only choose measurable strategies. Formally, $\psi_i : H \to [0, 1]$ is a measurable strategy for player i if, for all t, $\psi_i|_{H_t}$ is a measurable function from $(H_t, \mathcal{F}_t)$ to $([0, 1], \mathcal{B}([0, 1]))$, where $\mathcal{F}_t$ is the σ-algebra generated by the random vector $h_t : \Omega \to [-M, M]^t$ and $\mathcal{B}([0, 1])$ is the Borel σ-algebra on $[0, 1]$.

Definition 3. A strategy $\psi_i(h)$ for player i is *Markov* if $\psi_i(h) = \psi_i(h')$ for all $h = (\theta)^t$, $h' = (\theta)^{t'}$ such that $\theta_t(h) = \theta_{t'}(h')$.

A strategy ψ_1 for player 1 (2) is a *threshold strategy* with threshold θ^* if $\psi_1(h_t) = 1$ whenever $\theta_t(h_t) < \theta^*$ ($>$) and $\psi_1(h_t) = 0$ whenever $\theta_t(h_t) > \theta^*$ ($<$).

Definition 4. For any history h and strategies ψ_i, ψ_j, let $U_i(\psi_i, \psi_j|h)$ be i's continuation utility from strategy profile (ψ_i, ψ_j) at history h.

Definition 5. For any history h_0, any strategy ψ_i for player i, and any $\alpha \in [0, 1]$, let $\psi_i^{h_0,\alpha}$ be a strategy for player i given by $\psi_i^{h_0,\alpha}(h) = \alpha$ if $h = h_0$ and $\psi_i^{h_0,\alpha}(h) = \psi_i(h)$ otherwise.

Definition 6. Given two strategies $\psi_i(h)$, $\psi_i(h')$ for player i, we say $\psi_i \geq \psi_i'$ iff $\psi_i(h) \geq \psi_i'(h)$ for all histories h.

Proof of Proposition 1. Our proof of uniqueness uses familiar tools from the supermodular games literature (Milgrom and Roberts, 1990; Morris and Shin, 1998). The main steps are as follows. First, we show the game is supermodular in the following sense: if player i quits more in equilibrium, then player j's best response is to quit less, and vice versa. Standard results (Milgrom and Roberts, 1990) then imply the existence of a greatest equilibrium (where 1 quits as much as possible, and 2 as little as possible) and a smallest equilibrium (where the reverse happens) that serve as upper and lower bounds, respectively, for all other equilibria. Second, we show that the best response to a threshold strategy is a threshold strategy. Third, we note that the extremal equilibria can be found by iterating the best response correspondences from the extreme profiles where one player always surrenders and the other never does (which are degenerate threshold strategies), which implies that the extremal equilibria are in threshold strategies. Fourth, we show that there is a unique equilibrium in threshold strategies. Then the extremal equilibria must coincide. Since the extremal equilibria bound all other equilibria, there cannot be any other equilibria.

1. Supermodularity.

Definition 7. Given a strategy ψ_j for player j and a history h, let $V_i(\psi_j|h)$ be the highest continuation utility player i can attain conditional on the history being h and player j using strategy ψ_j, i.e.,

$$V_i(\psi_j|h) = \sup_{\psi_i} U_i(\psi_i, \psi_j|h).$$

Let $\tilde{V}_i(\psi_j|h)$ be the highest continuation utility player i can attain conditional on the history being h and player j using strategy ψ_j, if player i is restricted to not surrendering in the current period, i.e.,

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

Lemma 1. *Let $\psi_j \geq \psi'_j$ be two strategies for player j, let ψ_i be a strategy for player i and let h be any history. Then*

$$U_i(\psi_i, \psi_j|h) \leq U_i(\psi_i, \psi'_j|h).$$

Proof. Let h be a history for time t_0. Then $U_i(\psi_i, \psi_j|h) - U_i(\psi_i, \psi'_j|h) = \sum_{t=t_0}^{\infty} \delta^{t-t_0} E_t$, where E_t equals

$$\int Q(\psi_i, \psi'_j, (\theta)^t)(\psi_j((\theta)^t) - \psi'_j((\theta)^t))\psi_i((\theta^t)) \left(U_i\left(\psi_i^{(\theta)^t,1}, \psi_j^{(\theta)^t,1}|(\theta)^t\right) - H_1\right) dP((\theta)^t|(\theta)^{t_0}) \leq 0$$

Here $Q(\psi_i, \psi'_j, (\theta)^t)$ 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 ψ_i, ψ'_j respectively. The last inequality follows from the fact that $U_i\left(\psi_i^{(\theta)^t,1}, \psi_j^{(\theta)^t,1}|(\theta)^t\right) - H_1 < 0$ by Assumption B5. □

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

$$V_i(\psi_j|h) \leq V_i(\psi'_j|h).$$

Definition 8. We say a strategy ψ_i for player i is a *subgame-perfect response* to a strategy ψ_j for player j if it is a best response in every subgame.

Lemma 2. *ψ_i is a subgame-perfect response to ψ_j iff, for all h, it satisfies: $\psi_i(h) = 1$ if $\tilde{V}_i(\psi_j|h) > 0$ and $\psi_i(h) = 0$ if $\tilde{V}_i(\psi_j|h) < 0$.*

Proof. ($\Longrightarrow$) Starting at history h, any strategy ψ_i attains utility $U_i(\psi_i, \psi_j|h) = \psi_i(h)U(\psi_i^{h,1}, \psi_j|h)$. If $\tilde{V}_i(\psi_j|h) > 0$, attaining positive utility starting at h is possible for i, so a subgame-perfect response ψ_i must accomplish this. Then $U_i(\psi_i^{h,1}, \psi_j|h) > 0$, so it is optimal to set $\psi_i(h) = 1$. If $\tilde{V}_i(\psi_j|h) < 0$, then $U_i(\psi_i^{h,1}, \psi_j|h) < 0$ for any ψ_i, so it is optimal to set $\psi_i(h) = 0$.

($\Longleftarrow$) Suppose ψ_i satisfies the given condition but is not a subgame-perfect response, so there is ψ_i' such that $U_i(\psi_i, \psi_j|h) < U_i(\psi_i', \psi_j|h)$ for some h. The definition of ψ_i, plus ($\Longrightarrow$), imply that $\psi_i(h) = \psi_i'(h)$. Then there is some history h' in the next period such that $U_i(\psi_i', \psi_j|h') - U_i(\psi_i, \psi_j|h') > \frac{U_i(\psi_i', \psi_j|h) - U_i(\psi_i, \psi_j|h)}{\delta}$. Iterating yields a contradiction, since both ψ_i and ψ_i' can only generate payoffs in $[0, H_i]$ at every history. □

Lemma 3. *Let ψ_j, ψ_j' be two strategies for player j such that $\psi_j \geq \psi_j'$. Let $\psi_i \in BR_i(\psi_j)$. Then there is $\psi_i' \in BR_i(\psi_j')$ such that $\psi_i' \geq \psi_i$.*

Proof. From Corollary 1 and the definition of $\tilde{V}_i(\psi_j|h)$, we get $\tilde{V}_i(\psi_j|h) \leq \tilde{V}_i(\psi_j'|h)$ for all h.

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

Define ψ_i' as follows: $\psi_i'(h) = 1$ if $h \in A_+' \cup A_0'$ and $\psi_i'(h) = 0$ otherwise. Then $\psi_i' \geq \psi_i$ by construction, and ψ_i' is a best response to ψ_j' by Lemma 2. □

Lemma 3 implies the existence of extremal equilibria (Milgrom and Roberts, 1990). Since they turn out to be in threshold strategies, we develop some results on threshold strategies before characterizing the extremal equilibria.

2. Best response to threshold strategy is a threshold strategy.

Lemma 4. *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 payoff of 0, while continuing and surrendering tomorrow yields a strictly positive payoff. The proof for player 2 is identical. □

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 possibilities. Either player 2 surrenders at some time $t' \geq t$, or player 1 surrenders at some time $t' > t$. In the first 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 negative in this case. In the second case where player 1 surrenders, if the state never makes it below $-M_1$ before she surrenders, her utility is also negative. If the state makes it below $-M_1$, by B5, her payoff is no better than if would be if player 2 surrendered as soon as θ reached $-M_1$, and this payoff is negative by the same argument as in the first 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 $M - \eta$. By Lemma 1, player 1 would also prefer to surrender if player 2 used any other strategy that does not violate Lemma 4.

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. We denote player 1's best-response threshold by $T_1(\theta_*)$.*

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. We denote player 2's best-response threshold by $T_2(\theta^*)$.

Proof. We prove the first statement; the second one is analogous. Suppose that player 2 uses a threshold strategy with threshold $\theta_* \in [-M, M_2]$, which we denote by $\psi_2^{\theta_*}$. 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 $\psi_2^{\theta_*}$, i.e.,

$$V_1(\theta) = \sup_{\psi_1} U_1(\psi_1, \psi_2^{\theta_*}|\theta)$$

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

$$W(g)(\theta) = \begin{cases} H_1 & \text{if } \theta \leq \theta_* \\ \max\left(-c_1(\theta) + \delta E(g(\theta')|\theta), 0\right) & \text{if } \theta \in (\theta_*, M - \eta) \\ 0 & \text{if } \theta \in [M - \eta, M] \end{cases} \tag{8}$$

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 f_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 state $T_1(\theta_*) > \theta_*$ for which $\tilde{V}_1(T_1(\theta_*)) = 0$, or $\tilde{V}_1(\theta) < 0$ for all $\theta > \theta_*$.

By Lemma 2, in the first case, $\psi_1^{T_1(\theta_*)}$ is the essentially unique subgame-perfect response to $\psi_2^{\theta_*}$.[27] In the second case, the unique best response for player 1 is a threshold strategy with threshold $T_1(\theta_*) = \theta_*$, such that $\psi_1(\theta_*) = 1$. □

Corollary 2. T_1, $T_2 : [-M, M] \to [-M, M]$ *are weakly increasing.*

Proof. Follows from Lemma 3. □

3. Extremal equilibria exist and are in threshold strategies.

Lemma 7. *There is a greatest equilibrium and a smallest equilibrium, both in threshold strategies. That is, there are threshold strategy equilibria with thresholds $\underline{\theta}_*$ for player 2, $\underline{\theta}^*$ for player 1 and $\overline{\theta}_*$ for player 2, $\overline{\theta}^*$ for player 1 such that, for any SPE (ψ_1, ψ_2), $\psi_1^{\overline{\theta}^*} \leq \psi_1 \leq \psi_1^{\underline{\theta}^*}$ and $\psi_2^{\underline{\theta}_*} \leq \psi_2 \leq \psi_2^{\overline{\theta}_*}$.*

Proof. We use a standard argument similar to Milgrom and Roberts (1990). Define $\overline{\theta}_{*0} = -M$, $\overline{\theta}^*_0 = -M_1$ and $\overline{\theta}_{*(n+1)} = T_2(\overline{\theta}^*_n)$, $\overline{\theta}^*_{n+1} = T_1(\overline{\theta}_{*n})$ for all $n \geq 0$. Analogously, set $\underline{\theta}_{*0} = M_2$, $\underline{\theta}^*_0 = M$ and $\underline{\theta}_{*(n+1)} = T_2(\underline{\theta}^*_n)$, $\underline{\theta}^*_{n+1} = T_1(\underline{\theta}_{*n})$ for all $n \geq 0$.

Since $T_2(-M_1) \geq -M$ and $T_1(-M) \geq -M_1$, we have $\overline{\theta}_{*1} \geq \overline{\theta}_{*0}$, $\overline{\theta}^*_1 \geq \overline{\theta}^*_0$. Since T_1, T_2 are weakly increasing, iterating yields $\overline{\theta}_{*(n+1)} \geq \overline{\theta}_{*n}$, $\overline{\theta}^*_{n+1} \geq \overline{\theta}^*_n$ for all n. Let $\overline{\theta}_* = \lim_n \overline{\theta}_{*n}$, $\overline{\theta}^* = \lim_n \overline{\theta}^*_n$. Analogously $\underline{\theta}_{*(n+1)} \leq \underline{\theta}_{*n}$, $\underline{\theta}^*_{n+1} \leq \underline{\theta}^*_n$ for all n, and we set $\underline{\theta}_* = \lim_n \underline{\theta}_{*n}$, $\underline{\theta}^* = \lim_n \underline{\theta}^*_n$.

[27] It is not unique in the sense that any value $\psi_1(T_1(\theta_*)) \in [0, 1]$ is optimal.

Let (ψ_1, ψ_2) be an SPE. By construction, $\psi_1^{\overline{\theta}^*_0} \leq \psi_1 \leq \psi_1^{\underline{\theta}^*_0}$ and $\psi_2^{\underline{\theta}_{*0}} \leq \psi_2 \leq \psi_2^{\overline{\theta}_{*0}}$. Applying Lemma 3 and using that $\psi_2 \in BR(\psi_1)$, $\psi_1 \in BR(\psi_2)$, we obtain $\psi_1^{\overline{\theta}^*_1} \leq \psi_1 \leq \psi_1^{\underline{\theta}^*_1}$ and $\psi_2^{\underline{\theta}_{*1}} \leq \psi_2 \leq \psi_2^{\overline{\theta}_{*1}}$. Iterating, $\psi_1^{\overline{\theta}^*_n} \leq \psi_1 \leq \psi_1^{\underline{\theta}^*_n}$ and $\psi_2^{\underline{\theta}_{*n}} \leq \psi_2 \leq \psi_2^{\overline{\theta}_{*n}}$ for all n. Then $\psi_1^{\overline{\theta}^*} \leq \psi_1 \leq \psi_1^{\underline{\theta}^*}$ and $\psi_2^{\underline{\theta}_*} \leq \psi_2 \leq \psi_2^{\overline{\theta}_*}$ up to a set of measure zero. (We can tie-break within the threshold strategies so that the claim holds properly.)

It remains to show that $(\psi_1^{\underline{\theta}^*}, \psi_2^{\underline{\theta}_*})$ and $(\psi_1^{\overline{\theta}^*}, \psi_2^{\overline{\theta}_*})$ are equilibria. It is enough to show that $T_2(\underline{\theta}^*) = \underline{\theta}_*$, $T_1(\underline{\theta}_*) = \underline{\theta}^*$; the other one is analogous. Since $T_2(\underline{\theta}^*_n) = \underline{\theta}_{*(n+1)}$ and $T_1(\underline{\theta}_{*n}) = \underline{\theta}^*_{n+1}$ for all n, and $\underline{\theta}_* = \lim_n \underline{\theta}_{*n}$, $\underline{\theta}^* = \lim_n \underline{\theta}^*_n$, it is enough to show that T_1, T_2 are continuous. We show in the next step that, in fact, these functions are Lipschitz. □

4. Unique equilibrium in threshold strategies; extremal equilibria coincide.

By Lemma 6, if one player uses a threshold strategy, the other player does too, and the latter threshold is uniquely determined as a function of the former. An equilibrium in threshold strategies is given by a threshold θ^* for player 1 such that $T_1(T_2(\theta^*)) = \theta^*$.

We now argue that, for any $x > y$ such that $T_1(y) > y$, $T_1(x) - T_1(y) < x - y$. (In particular, T_1 is continuous.) 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_ϵ be the function $t_\epsilon(\theta) = \theta - \epsilon$. Take $\epsilon = \tilde{\theta} - \tilde{\theta}'$. For any function V, denote $\overline{V} = V \circ t_\epsilon$. For any operator W, define $\overline{W}$ by $\overline{W}(g) = W(g \circ t_\epsilon^{-1}) \circ t_\epsilon$.

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

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

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

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

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

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 $-c_1(\theta - \epsilon) + \delta E(f(\theta_{t+1} + \epsilon)|\theta_t = \theta - \epsilon) > 0$ for the same θ, whence $T_1(\tilde{\theta}') + \epsilon > T_1(\tilde{\theta})$. 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 $\theta^* \neq \theta^{*\prime}$ are both fixed points of $T_1 \circ T_2$, we would have that $|T_1(T_2(\theta^*)) - T_1(T_2(\theta^{*\prime}))| \leq |T_2(\theta^*) - T_2(\theta^{*\prime})| \leq |\theta^* - \theta^{*\prime}|$ with at least one strict inequality, a contradiction.

We showed in part 3 that $\overline{\theta}^*$, $\underline{\theta}^*$ are both fixed points of $T_1 \circ T_2$. Therefore $\overline{\theta}^* = \underline{\theta}^*$. Applying T_2, $\overline{\theta}_* = \underline{\theta}_*$ as well—the extremal equilibria coincide. Then Lemma 7 yields that all SPEs coincide with the unique threshold equilibrium, up to randomizing at the thresholds. □

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

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

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, which is to say 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 $\hat{\theta}^* \leq \theta^*$.

□

Proof of Proposition 3. We first characterize an equilibrium in threshold strategies (θ_*, θ^*) of the continuous time game, i.e., a fixed point of $T_1^0 \circ T_2^0$, which turns out to be unique (among threshold strategies). We then show that $\theta_*^\Delta \to \theta_*$, $\theta^{*\Delta} \to \theta^*$ as $\Delta \to 0$: (θ_*, θ^*) is also the unique limit of discrete-time equilibria.

Threshold-strategy equilibria in continuous time.

Begin by considering an arbitrary (not necessarily equilibrium) threshold strategy profile, with thresholds $\underline{\theta} < \overline{\theta}$. Note that both $V_i(\theta_t)$ and $P_i(\theta_t)$ are themselves drift-diffusion processes by Itô's lemma:

$$dV_i(\theta_t) = \left(\mu(\theta_t)V_i'(\theta_t) + \frac{\sigma^2}{2}V_i''(\theta_t)\right) dt + \sigma V_i(\theta_t) dB_t \tag{9}$$

$$dP_i(\theta_t) = \left(\mu(\theta_t)P_i'(\theta_t) + \frac{\sigma^2}{2}P_i''(\theta_t)\right) dt + \sigma P_i(\theta_t) dB_t \tag{10}$$

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

$$0 = [-c_i(\theta_t) - \gamma V_i(\theta_t)]dt + E(dV_i(\theta_t)),$$

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

Taking expectation of Equations 9 and 10 conditional on the value of θ_t,

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

In addition, the following boundary conditions must hold. $V_1(\underline{\theta}) = H_1$, $V_2(\underline{\theta}) = 0$ as 2 surrenders at this state; $V_1(\overline{\theta}) = 0$, , $V_2(\overline{\theta}) = H_2$ as 1 surrenders; and similarly $P_1(\underline{\theta}) = 1$, $P_1(\overline{\theta}) = 0$, $P_2(\underline{\theta}) = 0$, $P_2(\overline{\theta}) = 1$.

Define T_1^0 as follows: $\overline{\theta} = T_1^0(\underline{\theta})$ iff $\overline{\theta}$ is a best response threshold for player 1 to $\underline{\theta}$, and T_2^0 analogously. We now argue that $V_1'(\overline{\theta}) = 0$ iff $\overline{\theta} = T_1^0(\underline{\theta})$, and $V_2'(\underline{\theta}) = 0$ iff $\underline{\theta} = T_2^0(\overline{\theta})$—hence the smooth-pasting conditions pin down the equilibrium thresholds θ_*, θ^*. To see why, some machinery is required. For $\theta \in [\underline{\theta}, \overline{\theta}]$, let $Q_t(\theta)$ be the probability that, with the game having started in state θ, player 1 has not surrendered by time t; and define $Q(\theta) = \int_0^\infty e^{-\gamma t} Q_t(\theta)$. Define $\tilde{V}_1(\theta) = \frac{V_1(\theta)}{Q(\theta)}$ for $\theta < \overline{\theta}$, and $\tilde{V}_1(\overline{\theta}) = \lim_{\theta \to \overline{\theta}} \tilde{V}_1(\theta)$. Note that $\tilde{V}_1(\theta)$ is the expected utility of player 1 under the following assumptions: the initial state is θ; the stochastic process $(\theta_t)_t$ snaps back to θ if it ever hits $\overline{\theta}$; and player 1 never surrenders. In particular, $\tilde{V}_1(\overline{\theta})$ is the expected utility of player 1 under the following assumptions: the initial state is θ; the stochastic process $(\theta_t)_t$ is reflecting at $\overline{\theta}$; and player 1 never surrenders. It is clear that $\tilde{V}_1(\overline{\theta}) = 0$ $(>, <)$ iff $\overline{\theta} = T_1(\underline{\theta})$ $(<, >)$. In addition, since $V_1(\theta) = Q(\theta)\tilde{V}_1(\theta)$, for $\theta < \overline{\theta}$, we have

$$V_1'(\theta) = Q'(\theta)\tilde{V}_1(\theta) + Q(\theta)\tilde{V}_1'(\theta).$$

Taking the limit as $\theta \to \overline{\theta}$, we obtain

$$V_1'(\overline{\theta}) = Q'(\overline{\theta})\tilde{V}_1(\overline{\theta}) + Q(\overline{\theta})\tilde{V}_1'(\overline{\theta}) = Q'(\overline{\theta})\tilde{V}_1(\overline{\theta}),$$

as $Q(\overline{\theta}) = 0$. It can be shown that $Q'(\overline{\theta}) < 0$; the result follows.

Discrete-time equilibria converge to this equilibrium.

It can be shown by an analogous argument to Proposition 1 that the T_i^0 satisfy $|T_i^0(x) - T_i^0(y)| < |x - y|$. Thus the continuous time game has a unique equilibrium (θ_*, θ^*) among threshold strategy profiles, given by a fixed point of $T_1^0 \circ T_2^0$. We now argue that $\theta_*^\Delta \to \theta_*$, $\theta^{*\Delta} \to \theta^*$ as $\Delta \to 0$. Clearly, it is enough to show that $||T_i^\Delta - T_i^0||_\infty \to 0$ as $\Delta \to 0$ for $i = 1, 2$. In turn, because the T_i^Δ and T_i^0 are increasing and continuous on a compact interval, pointwise convergence implies uniform convergence, so it is enough to show that $T_i^\Delta(x) \to T_i^0(x)$ for any x.

We show this for $i = 1$; the other case is analogous. Note that, holding fixed player 2's threshold at x, player 1 is weakly worse off when $t = 0, \Delta, 2\Delta, \dots$ than in continuous time, for two reasons. First, as the state evolves continuously, player 1 would like to quit whenever θ_t crosses below $T_1^0(x)$, but in discrete time she must wait for the next time of the form $k\Delta$. Second, player 1 is always better off if player 2 quits, but player 2 quits less in discrete time, even for the same x because she also has to wait for times of the form $k\Delta$ to quit. Then $T_1^\Delta(x) \leq T_1^0(x)$. It is not hard to show that these losses vanish in the limit as $\Delta \to 0$, which yields $T_1^\Delta(x) \to T_1^0(x)$. □

Proof of Proposition 4. See Appendix B. □

Proof of Proposition 5. First, we will argue that $\theta^*(\nu) - \theta_*(\nu) \to 0$ as $\nu \to 0$. Suppose otherwise. Then, by the Bolzano-Weierstrass theorem, we can take a sequence $\nu_k \to 0$ such that $\theta_*(\nu_k) \to \theta_{**}$, $\theta^*(\nu_k) \to \theta^{**}$, and $\theta_{**} < \theta^{**}$.

Take the limit of the players' equilibrium value functions $V_i(\theta; \nu_k)$ as $k \to \infty$, and denote the limit by $V_i(\theta; 0)$. It follows that $V_1(\cdot; 0)$, $V_2(\cdot; 0)$ must satisfy the equilibrium conditions in Proposition 3, in particular Equation 1, over the disputed region $[\theta_{**}, \theta^{**}]$, with the parameters $\mu(\theta) \equiv 0$, $\sigma^2 = 0$, i.e., they must satisfy $c(\theta) + V_i(\theta) \equiv 0$, which is impossible. Hence $\theta^*(\nu) - \theta_*(\nu) \to 0$ as $\nu \to 0$.

At this point a digression is needed. For any fixed $\overline{\theta} \in [-M, M]$, consider a degenerate version of the model in which $\mu(\theta) \equiv \mu(\overline{\theta})$, $c_1(\theta) \equiv c_1(\overline{\theta})$ and $c_2(\theta) \equiv c_2(\overline{\theta})$ for all θ. Clearly in this model the mapping $T_1(T_2(\theta))$ is of the form $\theta + \Delta$ for some Δ. Say $\overline{\theta}$ is 1-favored if $\Delta > 0$, 2-favored if $\Delta < 0$, and balanced if $\Delta = 0$. Note that Δ is a continuous and strictly decreasing function of $\overline{\theta}$, and so there is a unique balanced state. We denote this state by θ^l.

Now the crucial observation is that, for any value of ν, we must have $\theta^l \in [\theta_*(\nu), \theta^*(\nu)]$. Indeed, if this were not the case, then the interval $[\theta_*(\nu), \theta^*(\nu)]$ would be made up entirely of 1-favored states or entirely of 2-favored states, in which case we would have $T_1(T_2(\theta^*(\nu)) > \theta^*(\nu)$ or $T_1(T_2(\theta^*(\nu)) < \theta^*(\nu)$ respectively, a contradiction. Then, as $\theta^*(\nu) - \theta_*(\nu) \to 0$, it must be that $\theta_*(\nu), \theta^*(\nu) \to \theta^l$.

Finally, if $\mu \equiv 0$, then clearly the only balanced state θ^l is the one with $c_1(\theta^l) = c_2(\theta^l) =: c$. To characterize welfare under this condition, let us first calculate V_1, V_2 in a degenerate case in which $c_1(\theta)$, $c_2(\theta) \equiv c$, and imposing by fiat that $\theta_* + \theta^* = 0$ (since this game otherwise has multiple equilibria.) Equation 1 becomes $c^* + \gamma V_i(\theta) = \frac{\sigma^2}{2} V_i''(\theta)$, which together with the boundary and smooth-pasting conditions from Proposition 3 yields the solution

$$V_1(\theta) = -\frac{c}{\gamma} + \frac{c}{2\gamma}e^{\frac{\sqrt{2\gamma}}{\sigma}(\theta^*-\theta)} + \frac{c}{2\gamma}e^{-\frac{\sqrt{2\gamma}}{\sigma}(\theta^*-\theta)}$$
$$V_2(\theta) = -\frac{c}{\gamma} + \frac{c}{2\gamma}e^{\frac{\sqrt{2\gamma}}{\sigma}(\theta^*+\theta)} + \frac{c}{2\gamma}e^{-\frac{\sqrt{2\gamma}}{\sigma}(\theta^*+\theta)}$$

for $\theta \in [-\theta^*, \theta^*]$, and θ^* must be such that

$$H + \frac{c}{\gamma} = \frac{c}{2\gamma}e^{2\frac{\sqrt{2\gamma}}{\sigma}\theta^*} + \frac{c}{2\gamma}e^{-2\frac{\sqrt{2\gamma}}{\sigma}.\theta^*}$$

Denote $Y = e^{2\frac{\sqrt{2\gamma}}{\sigma}\theta^*}$ and $X = \sqrt{Y}$. Then this is equivalent to

$$Y + \frac{1}{Y} = \frac{2\gamma H}{c} + 2$$
$$Y^2 - \left(\frac{2\gamma H}{c} + 2\right)Y + 1 = 0 \Longrightarrow Y = \frac{\gamma H}{c} + 1 + \sqrt{\left(\frac{\gamma H}{c} + 1\right)^2 - 1}.$$

(This is the only valid solution since $Y \geq 1$ by construction.) Then

$$V_i\left(\frac{\theta_* + \theta^*}{2}\right) = V_i(0) = -\frac{c}{\gamma} + \frac{c}{2\gamma}e^{\frac{\sqrt{2\gamma}}{\sigma}\theta^*} + \frac{c}{2\gamma}e^{-\frac{\sqrt{2\gamma}}{\sigma}\theta^*} = \frac{c}{2\gamma}\left(X + \frac{1}{X} - 2\right) = \frac{c}{2\gamma}\frac{(X-1)^2}{X}.$$

Since $\frac{(X^2-1)^2}{X^2} = X^2 + \frac{1}{X^2} - 2 = \frac{2\gamma H}{c}$ by construction,

$$\frac{c}{2\gamma}\frac{(X-1)^2}{X} = \frac{c}{2\gamma}\frac{(X-1)^2}{X}\frac{(X+1)^2}{X}\frac{X}{(X+1)^2} = H\frac{1}{X + \frac{1}{X} + 2},$$

as we wanted.

The general result now follows from a continuity argument. In this degenerate case where $c_i(\theta) \equiv c$, the result $V_i\left(\frac{\theta_*+\theta^*}{2}\right) = H\frac{1}{X+\frac{1}{X}+2}$ holds exactly for any σ. The setting where c_1, c_2 are general cost functions satisfying B1-5 plus $c_1(0) = c_2(0)$ and $\tilde{\sigma} = \sqrt{\nu}\sigma$ for ν small is equivalent, through the change of variables $\tilde{\theta} = \frac{\theta}{\sqrt{\nu}}$, to one where the noise parameter is held constant at σ and we redefine $\tilde{c}_i(x) = c_i(\sqrt{\nu}x)$. This model converges to the degenerate case as $\nu \to 0$. Since the solution to the ODE system in Proposition 3 is continuous in the c_i, $V_i\left(\frac{\theta_*(\nu)+\theta^*(\nu)}{2}\right)$ must go to the limit value we characterized as $\nu \to 0$. □

Proof of Proposition 6. Let $V_i(\theta; x)$ be the players' continuation values after player 1 makes a one-time concession which cedes $[x, 1]$. As we have assumed $\mu \equiv 0$ and $\gamma = 0$, the continuation values can be calculated from Equations 3–4, taking $\theta_* = \theta_*(x)$, $\theta^* = \theta^*(x)$.

More precisely,

$$V_1(\theta;x) = \min\left(\frac{2}{\sigma^2}\int_{\theta}^{\theta^*(x)}(\lambda-\theta)c_1(\lambda)d\lambda, \int_0^x v_1(\tilde{x})d\tilde{x}\right).$$

That is, $V_1(\theta;x) = \frac{2}{\sigma^2}\int_{\theta}^{\theta^*(x)}(\lambda-\theta)c_1(\lambda)d\lambda$ if $\theta \geq \theta_*(x)$ and $V_1(\theta;x) = \int_0^x v_1(\tilde{x})d\tilde{x}$ otherwise. Note that the first expression depends only on x through $\theta^*(x)$ and is an increasing function of $\theta^*(x)$, while the second expression is increasing in x.

From Proposition 4.(i)-(iii), we know that $\theta_*(x)$ increases in the size of the concession (i.e., it is decreasing in x), since a marginal increase in player 1's concession always lowers player 2's remaining prize value proportionally more than her own. Thus, there is $x_0 \geq 0$ such that $\theta_*(x) > \theta$ iff $x < x_0$. Choosing $x^* = x_0$ then dominates any choice below x_0 (which would still induce immediate surrender but leave a smaller prize for player 1). Player 1 thus simply maximizes $\theta^*(x)$ over $x \in [x_0, 1]$—an equivalent condition to $\theta_*(x) \leq \theta$.

Since the continuation game after any concession x has a unique equilibrium (Proposition 3), to show equilibrium existence, it is enough to show that player 1's optimization problem, $\max_{x\in[x_0,1]} \theta^*(x)$, has a solution. Define $\theta_*(H_1, H_2)$, $\theta^*(H_1, H_2)$ as the equilibrium thresholds of the baseline model with generic prizes H_1, H_2, i.e., as the solutions of the system

$$H_1 = \frac{2}{\sigma^2}\int_{\theta_*}^{\theta^*}(\lambda-\theta_*)c_1(\lambda)d\lambda, \qquad H_2 = \frac{2}{\sigma^2}\int_{\theta_*}^{\theta^*}(\theta^*-\lambda)c_2(\lambda)d\lambda.$$

By Proposition 3, $\theta_*(H_1, H_2)$, $\theta^*(H_1, H_2)$ exist and are unique for all $H_1, H_2 > 0$. Moreover, by the inverse function theorem, they are C^1 functions of (H_1, H_2), since the mapping $(\theta_*, \theta^*) \mapsto (H_1, H_2)$ has Jacobian

$$\frac{2}{\sigma^2}\begin{pmatrix} -\int_{\theta_*}^{\theta^*} c_1(\lambda)d\lambda & (\theta^*-\theta_*)c_1(\theta^*) \\ -(\theta^*-\theta_*)c_2(\theta_*) & \int_{\theta_*}^{\theta^*} c_2(\lambda)d\lambda \end{pmatrix}$$

which is positive as $(\theta^*-\theta_*)c_1(\theta^*) > \int_{\theta_*}^{\theta^*} c_1(\lambda)d\lambda$ and $(\theta^*-\theta_*)c_2(\theta_*) > \int_{\theta_*}^{\theta^*} c_2(\lambda)d\lambda$. Then $(\theta_*(x), \theta^*(x)) = (\theta_*(\tilde{H}_1(x), \tilde{H}_2(x)), \theta^*(\tilde{H}_1(x), \tilde{H}_2(x))$ is continuous because $(\tilde{H}_1, \tilde{H}_2)$ is continuous in x. In particular, $x \mapsto \theta^*(x)$ is continuous, so it attains a maximum in $[x_0, 1]$. □

Proof of Proposition 7. For the first part, consider the following strategy profile: player 1 makes the optimal one-shot concession at $t = 0$, and afterwards player 1 (2) surrenders if θ_t hits $\theta^*(x^*)$ ($\theta_*(x^*)$). If player 1 deviates by not conceding or under-conceding at $t = 0$, player 2 expects player 1 to concede down to $[0, x^*)$ immediately. If player 1 deviates by over-conceding, the two players expect a continuation in which player 1 concedes to the optimal concession below x^* and then no further partial concessions take place, and player 2 best-responds, etc.

Given this strategy profile, player 2 has no incentive to surrender at $t = 0$. Afterwards, since player 1 plays as in Proposition 6, player 2's best response is to surrender whenever $\theta_t \leq \theta_*(x^*)$. By construction, player 1's payoff is strictly lower if she over-concedes, and no better if she under-concedes (indeed, her payoff will be strictly lower if she fails to concede down to $[0, x^*)$ by the time θ_t first reaches $\{\theta_*(x^*), \theta^*(x^*)\}$).

We prove the second part by induction on k. For $k = 1$, we are in the baseline game. For $k = 2$, assume an equilibrium in threshold strategies in which player 2 surrenders when $\theta_t \leq \theta_*(0, 1)$ and player 1 concedes down to $[0, x_1)$ when $\theta_t \geq \theta^*(0, 1)$. Of course, the equilibrium after a concession is pinned down by Proposition 1, with $\theta_*(0, x_1) = \tilde{\theta}_*(0, x_1)$, $\theta^*(0, x_1) = \tilde{\theta}^*(0, x_1)$.

There are two cases: $x^* = 1$ or $x^* = x_1$. If $x^* = 1$, we want to show that player 1's equilibrium payoff is no higher than $V_1(\theta_0)$, where $V_1(\theta) = V_1(\theta; 0, 1)$ is defined as in the baseline model. By similar arguments as in Proposition 3, we can show that player 1's best response as a function of $\theta_*(0, 1)$ satisfies the following: if $\theta_*(0, 1) < T_1^{-1}(\tilde{\theta}^*(0, x_1))$, then it is optimal for player 1 to concede down to $[0, x_1)$ for $\theta_t \geq \tilde{\theta}_*(0, x_1)$, so $\theta^*(0, 1) \leq \tilde{\theta}_*(0, x_1)$. (Again, the T_i are the same functions as in the baseline model, and the $\tilde{\theta}(\cdot)$ are defined as in the two-sided concessions model.) In this case, player 1's ex ante payoff is exactly $V_1(\theta_0; 0, x_1)$, her continuation payoff after a concession, which by the assumption that $x^* = 1$ and (3) is lower than $V_1(\theta_0)$. If $\theta_*(0, 1) = T_1^{-1}(\tilde{\theta}^*(0, x_1))$, player 1 is indifferent about conceding for $\theta_t \in (\tilde{\theta}_*(0, x_1), \tilde{\theta}^*(0, x_1))$ and her ex ante payoff is $V_1(\theta_0)$ whatever she does.

Finally, if $\theta_*(0, 1) > T_1^{-1}(\tilde{\theta}^*(0, x_1))$, then it is optimal for player 1 to concede down to $[0, x_1)$ only for $\theta_t \geq T_1(\theta_*(0, 1))$, so $\theta^*(0, 1) = T_1(\theta_*(0, 1)) > \tilde{\theta}^*(0, x_1)$ by construction, i.e., a partial concession must lead to immediate surrender by player 1. But then the conditions pinning down $\theta_*(0, 1)$ and $\theta^*(0, 1)$ are the same as in the model without concessions, whence $\theta^*(0, 1) = \tilde{\theta}^*(0, 1)$ and player 1's ex ante payoff is exactly $V_1(\theta_0)$ as in 6.

If $x^* = x_1$, a similar argument applies. Briefly, if $\theta_*(0, 1) \leq T_1^{-1}(\tilde{\theta}^*(0, x_1))$, then player 1 is at least weakly willing to concede at $t = 0$ and her ex ante payoff is as in Proposition 6. If $\theta_*(0, 1) > T_1^{-1}(\tilde{\theta}^*(0, x_1))$, then $\theta^*(0, 1) = T_1(\theta_*(0, 1)) > \tilde{\theta}^*(0, x_1)$, so a concession by player 1 leads to her surrender; the thresholds are as in the model without concessions; and player 1's payoff would then have to be $V_1(\theta_0)$, which by the assumption that $x^* = x_1$ is lower than $V_1(\theta_0; 0, x_1)$.

Now suppose the result is true for a general k_0, and take $k = k_0 + 1$. Denote by $\tilde{V}_1(\theta)$ player 1's continuation payoff after a concession to $[0, x_{k-1})$. If this payoff is as would be obtained in the equilibrium of Proposition 6 (that is, $\tilde{V}_1(\theta) = V_1(\theta; 0, \hat{x})$, where $\hat{x}$ is the optimal concession subject to the restriction $x_i \leq x_{k-1}$) then the same argument applies as in the case $k = 2$; indeed, $\hat{x}$ plays the role of x_1 in the preceding argument. If player 1's payoff is any lower, that is, $\tilde{V}_1(\theta) \leq V_1(\theta; 0, \hat{x})$ for all θ, then the same argument as in the case $k = 2$ now implies that her ex ante payoff is at most

$$\max(\tilde{V}_1(\theta), V_1(\theta)) \leq \max(V_1(\theta; 0, \hat{x}), V_1(\theta)) = V_1(\theta; 0, x^*).$$

Finally, $\tilde{V}_1(\theta)$ cannot be higher than $V_1(\theta; 0, \hat{x})$ by the inductive hypothesis.

□

Proof of Remark 1. Since $(H_1'(x), H_2'(x)) = (v_1(x), v_2(x))$, the slope of the feasible path of prize pairs is $\frac{v_2(x)}{v_1(x)}$. This is increasing in x since $v_2(x)$ is increasing and $v_1(x)$ is decreasing, so the feasible path of prize pairs is strictly convex.

As for the level curve $H_2(H_1; \theta^*)$, note that, using (3)-(4),

$$\frac{\partial H_2(H_1; \theta^*)}{\partial H_1} = \frac{\frac{\partial H_2(\theta_*, \theta^*)}{\partial \theta_*}}{\frac{\partial H_1(\theta_*, \theta^*)}{\partial \theta_*}} = \frac{\frac{\partial}{\partial \theta_*}\left(\frac{2}{\sigma^2}\int_{\theta_*}^{\theta^*}(\theta^* - \lambda)c_2(\lambda)d\lambda\right)}{\frac{\partial}{\partial \theta_*}\left(\frac{2}{\sigma^2}\int_{\theta_*}^{\theta^*}(\lambda - \theta_*)c_1(\lambda)d\lambda\right)} =$$
$$= \frac{-(\theta^* - \theta_*)c_2(\theta_*)}{\int_{\theta_*}^{\theta^*} -c_1(\lambda)d\lambda} = \frac{c_2(\theta_*)}{⨍_{\theta_*}^{\theta^*} c_1(\lambda)d\lambda}.$$

As H_1 grows while keeping θ^* constant, θ_* decreases by Proposition 4. Then, by B1, $c_2(\theta_*)$ increases and $⨍_{\theta_*}^{\theta^*} c_1(\lambda)d\lambda$ decreases, so $\frac{\partial H_2}{\partial H_1}$ increases. Thus the level curves are strictly convex. □

B Additional Proofs (For Online Publication)

Denote by $H_1(\theta_*, \theta^*)$ the value of H_1 which makes θ^* the optimal surrender threshold for player 1 when player 2's threshold is θ_*. Define $H_2(\theta_*, \theta^*)$ analogously.

Lemma 8. $\frac{|\frac{\partial H_1(\theta_*,\theta^*)}{\partial \theta_*}|}{|\frac{\partial H_2(\theta_*,\theta^*)}{\partial \theta_*}|}$ *is increasing in* θ_* *for all* $\theta_* < \theta^*$.

Proof. We proceed in several steps. The general strategy of the proof will be to identify values of the parameters (in particular c_1, c_2 and μ) that yield the tightest case, in the sense that $|\frac{\partial H_1(\theta_*,\theta^*)}{\partial \theta_*}|$ decreases as fast as possible and $|\frac{\partial H_2(\theta_*,\theta^*)}{\partial \theta_*}|$ increases as fast as possible, and then prove the result directly in that case.

We begin with $|\frac{\partial H_2(\theta_*,\theta^*)}{\partial \theta_*}|$. First, some auxiliary definitions. Define $V_1(\theta; \underline{\theta}, \overline{\theta}, c_1, H)$ as player 1's expected payoff under the following conditions: the initial state is θ, the (possibly non-equilibrium) disputed region is $[\underline{\theta}, \overline{\theta}]$, and player 1's cost function and prizes are c_1, H. Define $\hat{V}_1(\theta; \underline{\theta}, \overline{\theta}, c_1, H)$ as the same object but under the additional assumption that the stochastic process $(\theta_t)_t$ is reflecting at $\overline{\theta}$, and player 1 never surrenders. (Note that $\hat{V}_1(\theta; \underline{\theta}, \overline{\theta}, c_1, H) = V_1(\theta; \underline{\theta}, \overline{\theta}, c_1, H)$ whenever $\overline{\theta} = T_1(\underline{\theta})$.) Define V_2, $\hat{V}_2$ analogously.

Normalize $\gamma = 1$. We can write

$$\hat{V}_2(\theta_*; \theta_*, \theta^*, c_2, H) = -\int_{\theta_*}^{\theta^*} p(\theta) c_2(\theta) d\theta + p(\theta^*) H,$$

where $p(\theta)$ are probability weights satisfying $p(\theta^*) = 1 - \int_{\theta_*}^{\theta^*} p(\theta) d\theta$ and independent of c_2, H. In addition, note that

$$0 = V_2(\theta_*; \theta_*, \theta^*, c_2, H_2(\theta_*, \theta^*)) = \hat{V}_2(\theta_*; \theta_*, \theta^*, c_2, H_2(\theta_*, \theta^*)).$$

Now, for any $\epsilon > 0$, write

$$\begin{aligned} 0 &= V_2(\theta_* - \epsilon; \theta_* - \epsilon, \theta^*, c_2, H_2(\theta_* - \epsilon, \theta^*)) = \hat{V}_2(\theta_* - \epsilon; \theta_* - \epsilon, \theta^*, c_2, H_2(\theta_* - \epsilon, \theta^*)) = \\ &= -\int_{\theta_*-\epsilon}^{\theta_*} p(\theta) c_2(\theta) d\theta + Q\hat{V}_2(\theta_*; \theta_*, \theta^*, c_2, H_2(\theta_* - \epsilon, \theta^*)) = \\ &= -\int_{\theta_*-\epsilon}^{\theta_*} p(\theta) c_2(\theta) d\theta + Q\left[-\int_{\theta_*}^{\theta^*} p(\theta) c_2(\theta) d\theta + p(\theta^*) H_2(\theta_* - \epsilon, \theta^*)\right], \end{aligned}$$

where $p(\theta)|_{[\theta_*,\theta^*]}$ is the same function as in the previous expression; $p(\theta)_{[\theta_*-\epsilon,\theta_*]}$ is the probability of the future state being θ (before or after hitting θ_*); and $Q = 1 - \int_{\theta_*-\epsilon}^{\theta_*} p(\theta) d\theta$.

Equivalently, we can write

$$0 = -\int_{\theta_*-\epsilon}^{\theta_*} p(\theta)c_2(\theta)d\theta + \left(1 - \int_{\theta_*-\epsilon}^{\theta_*} p(\theta)d\theta\right) p(\theta^*) \left[H_2(\theta_* - \epsilon, \theta^*) - H_2(\theta_*, \theta^*)\right].$$

Taking $\epsilon \to 0$, we obtain

$$|\frac{\partial H_2(\theta_*, \theta^*)}{\partial \theta_*}| = -\frac{\partial H_2(\theta_*, \theta^*)}{\partial \theta_*} = \frac{p(\theta_*)c_2(\theta_*)}{p(\theta^*)}. \tag{11}$$

We are interested in the case where this expression increases as fast as possible, in the following sense. Take parameters c_2, μ and $\tilde{c}_2$, $\tilde{\mu}$, and denote by H_2, $\tilde{H}_2$ the corresponding prizes. We say $(\tilde{c}_2, \tilde{\mu}) > (c_2, \mu)$ if $\frac{|\frac{\partial \tilde{H}_2(\theta_*,\theta^*)}{\partial \theta_*}|}{|\frac{\partial H_2(\theta_*,\theta^*)}{\partial \theta_*}|}$ is increasing in θ_*. Then, if we prove Lemma 8 for $(\tilde{c}_2, \tilde{\mu})$, we have automatically proved it for (c_2, μ).

Suppose c_2 is strictly decreasing, and let $\tilde{c}_2$ be a constant positive cost. By Equation 11, we clearly have $(\tilde{c}_2, \mu) > (c_2, \mu)$. We can then restrict our attention to the case where c_2 is constant. WLOG, we take $c_2 \equiv 1$.

Next we will consider changes in μ. Write

$$|\frac{\partial H_2(\theta_*, \theta^*)}{\partial \theta_*}| = \frac{p(\theta_*; \theta_*, [\theta_*, \theta^*])}{p(\theta^*; \theta_*, [\theta_*, \theta^*])} := A(\theta_*),$$

where we have made it explicit that $p(\theta; \theta_0, [\theta_*, \theta^*])$ depends on the initial state, θ_0, and the disputed region, $[\theta_*, \theta^*]$. Take $\epsilon > 0$, and write

$$\frac{A(\theta_* + \epsilon)}{A(\theta_*)} = \frac{p(\theta_* + \epsilon; \theta_* + \epsilon, [\theta_* + \epsilon, \theta^*])}{p(\theta_*; \theta_*[\theta_*, \theta^*])} \frac{p(\theta^*; \theta_*, [\theta_*, \theta^*])}{p(\theta^*; \theta_* + \epsilon, [\theta_* + \epsilon, \theta^*])}$$

Now note that

$$p(\theta^*; \theta_*, [\theta_*, \theta^*]) = Qp(\theta^*; \theta_* + \epsilon, [\theta_* + \epsilon, \theta^*])$$
$$p(\theta_* + \epsilon; \theta_*, [\theta_*, \theta^*]) = Qp(\theta_* + \epsilon; \theta_* + \epsilon, [\theta_* + \epsilon, \theta^*]),$$

where $Q = \left(1 - \int_{\theta_*}^{\theta_*+\epsilon} p(\theta; \theta_*, [\theta_*, \theta^*])d\theta\right)$. Hence

$$\frac{A(\theta_* + \epsilon)}{A(\theta_*)} = \frac{p(\theta_* + \epsilon; \theta_*, [\theta_*, \theta^*])}{p(\theta_*; \theta_*, [\theta_*, \theta^*])}.$$

To rewrite this expression in a useful way we will need the following construction. Assume that $(\theta_t)_t$ starts at θ_* and follows the usual drift-diffusion process, but that $(\theta_t)_t$ is now reflecting at θ_* **and** absorbing at $\theta_* + \epsilon$. For $\theta \in [\theta_*, \theta^+\epsilon)$, define $\tilde{p}(\theta)$ as the probability that

$\theta_t = \theta$, and define $\tilde{p}(\theta_* + \epsilon) = 1 - \int_{\theta_*}^{\theta+\epsilon} \tilde{p}(\theta)d\theta$. Then, for $\theta \in [\theta_*, \theta_* + \epsilon]$,

$$p(\theta; \theta_*, [\theta_*, \theta^*]) = \tilde{p}(\theta) + \tilde{p}(\theta_* + \epsilon)p(\theta; \theta_* + \epsilon, [\theta_*, \theta^*])$$
$$\implies \frac{A(\theta_* + \epsilon)}{A(\theta_*)} = \frac{\tilde{p}(\theta_* + \epsilon)p(\theta_* + \epsilon; \theta_* + \epsilon, [\theta_*, \theta^*])}{\tilde{p}(\theta_*) + \tilde{p}(\theta_* + \epsilon)p(\theta_*; \theta_* + \epsilon, [\theta_*, \theta^*])}.$$

Now consider alternative non-decreasing drift functions $\tilde{\mu}$ such that $\tilde{\mu}(\theta) = \mu(\theta)$ for $\theta \in [\theta_*, \theta_* + \epsilon]$. Note that decreasing $\tilde{\mu}|_{[\theta_*+\epsilon,\theta^*]}$ increases $p(\theta_* + \epsilon; \theta_* + \epsilon, [\theta_*, \theta^*])$ and $p(\theta_*; \theta_* + \epsilon, [\theta_*, \theta^*])$ proportionally, without affecting $\tilde{p}(\theta_* + \epsilon)$ or $\tilde{p}(\theta_*)$, and hence it increases $\frac{A(\theta_*+\epsilon)}{A(\theta_*)}$. Hence, in order to maximize $\frac{A(\theta_*+\epsilon)}{A(\theta_*)}$, it is optimal to take $\tilde{\mu}(\theta) = \mu(\theta_*+\epsilon)$ for all $\theta \in [\theta_*+\epsilon, \theta^*]$. Taking the limit as $\epsilon \to 0$, it follows that in order to maximize $|\frac{\partial H_2(\theta_*,\theta^*)}{\partial \theta_*}|$ at a certain value of θ_*, it is optimal to take μ constant over $[\theta_*, \theta^*]$. (For now, the optimal μ might be a function of θ_*.)

Next we work with $|\frac{\partial H_1(\theta_*,\theta^*)}{\partial \theta_*}|$. Take $\epsilon > 0$ and write

$$\hat{V}_1(\theta_*; \theta_*, \theta^*, c_1, H_1(\theta_*, \theta^*)) = \hat{V}_1(\theta_*; \theta_* - \epsilon, \theta^*, c_1, H_1(\theta_* - \epsilon, \theta^*))$$
$$H_1(\theta_*, \theta^*) = -\int_{\theta_*-\epsilon}^{\theta^*} p(\theta)c_1(\theta)d\theta + \hat{p}(\theta_* - \epsilon)H_1(\theta_* - \epsilon, \theta^*),$$

where $p(\theta) = p(\theta; \theta_*, [\theta_* - \epsilon, \theta^*])$ and the process is assumed to be reflecting at θ^*, and $\hat{p}(\theta_* - \epsilon) = 1 - \int_{\theta_*-\epsilon}^{\theta^*} p(\theta)d\theta$. Rearranging,

$$H_1(\theta_* - \epsilon, \theta^*) - H_1(\theta_*, \theta^*) = \int_{\theta_*-\epsilon}^{\theta^*} p(\theta)(c_1(\theta) + H_1(\theta_* - \epsilon, \theta^*))d\theta$$

Let $\tilde{p}(\theta; \theta_*, [\theta_*, \theta^*])$ be the probability of θ_t being equal to θ in the future, when the initial state is θ_*, the disputed region is $[\theta_*, \theta^*]$, and the stochastic process governing $(\theta_t)_t$ is reflecting at θ_* and θ^*. Then, for $\theta \geq \theta_*$, $p(\theta) = \left(1 - \hat{p}(\theta_* - \epsilon) - \int_{\theta_*-\epsilon}^{\theta_*} p(\theta)d\theta\right) \tilde{p}(\theta; \theta_*, [\theta_*, \theta^*])$. It can be shown that $\int_{\theta_*-\epsilon}^{\theta_*} p(\theta)d\theta \in \mathcal{O}(\epsilon^2)$ for small ϵ, and of course $\hat{p}(\theta_*) = 1$. So, taking $\epsilon \to 0$,

$$|\frac{\partial H_1(\theta_*, \theta^*)}{\partial \theta_*}| = |\hat{p}'(\theta_*)| \int_{\theta_*}^{\theta^*} \tilde{p}(\theta; \theta_*, [\theta_*, \theta^*]))(c_1(\theta) + H_1(\theta_*, \theta^*))d\theta.$$

($\hat{p}'(\theta_*)$ is only a left-derivative as $\hat{p}(\theta)$ is undefined for $\theta > \theta_*$.)

Now note that, for any $\epsilon > 0$, there is a fixed $K \in (0, 1)$ such that $\tilde{p}(\theta; \theta_*, [\theta_*, \theta^*]) = K\tilde{p}(\theta; \theta_* + \epsilon, [\theta_* + \epsilon, \theta^*])$ for all $\theta \in [\theta_* + \epsilon, \theta^*]$. Hence, denoting $\tilde{p}(\theta) = \tilde{p}(\theta; \theta_*, [\theta_*, \theta^*])$,

$$\frac{|\frac{\partial H_1(\theta_*+\epsilon,\theta^*)}{\partial \theta_*}|}{|\frac{\partial H_1(\theta_*,\theta^*)}{\partial \theta_*}|} = \frac{|\hat{p}'(\theta_* + \epsilon; \theta_* + \epsilon)|}{|\hat{p}'(\theta_*; \theta_*)|} \frac{\int_{\theta_*+\epsilon}^{\theta^*} \tilde{p}(\theta)(c_1(\theta) + H_1(\theta_* + \epsilon, \theta^*))d\theta}{\int_{\theta_*}^{\theta^*} \tilde{p}(\theta)(c_1(\theta) + H_1(\theta_*, \theta^*))d\theta} \frac{\int_{\theta_*}^{\theta^*} \tilde{p}(\theta)d\theta}{\int_{\theta_*+\epsilon}^{\theta^*} \tilde{p}(\theta)d\theta}. \quad (12)$$

Note that by construction $\int_{\theta_*}^{\theta^*} \tilde{p}(\theta)d\theta = 1$. Then second factor in Equation 12 is approximately

$$
\begin{aligned}
&1 - \epsilon \frac{\tilde{p}(\theta_*)(c_1(\theta_*) + H_1(\theta_*, \theta^*)) + \int_{\theta_*}^{\theta^*} \tilde{p}(\theta)|\frac{\partial H_1(\theta_*,\theta^*)}{\partial \theta_*}|}{\int_{\theta_*}^{\theta^*} \tilde{p}(\theta)(c_1(\theta) + H_1(\theta_*, \theta^*))d\theta} \\
=&1 - \epsilon \frac{\tilde{p}(\theta_*)(c_1(\theta_*) + H_1(\theta_*, \theta^*)) + |\frac{\partial H_1(\theta_*,\theta^*)}{\partial \theta_*}|}{\int_{\theta_*}^{\theta^*} \tilde{p}(\theta)(c_1(\theta) + H_1(\theta_*, \theta^*))d\theta} \\
=&1 - \epsilon \frac{\tilde{p}(\theta_*)(c_1(\theta_*) + H_1(\theta_*, \theta^*)) + |\hat{p}'(\theta_*)| \int_{\theta_*}^{\theta^*} \tilde{p}(\theta; \theta_*, [\theta_*, \theta^*]))(c_1(\theta) + H_1(\theta_*, \theta^*))d\theta}{\int_{\theta_*}^{\theta^*} \tilde{p}(\theta)(c_1(\theta) + H_1(\theta_*, \theta^*))d\theta} \\
=&1 - \epsilon \frac{\tilde{p}(\theta_*)(c_1(\theta_*) + H_1(\theta_*, \theta^*))}{\int_{\theta_*}^{\theta^*} \tilde{p}(\theta)(c_1(\theta) + H_1(\theta_*, \theta^*))d\theta} - \epsilon|\hat{p}'(\theta_*)|.
\end{aligned}
$$

This expression equals $1 - \epsilon\tilde{p}(\theta_*) - \epsilon|\hat{p}'(\theta_*)|$ if c_1 is constant, and is strictly larger otherwise. Since c_1 appears nowhere else in Equation 12, it follows that in order to make the RHS of Equation 12 as small as possible it is optimal to take c_1 constant. Next, we will argue briefly that it is optimal to also take μ constant. Denoting $H_1(\theta_*, \theta^*) = V_1(\theta_*)$, note that we are effectively trying to minimize $\frac{V_1''(\theta_*)}{V_1'(\theta_*)}$. Recall that $V_1(\theta)$ must solve Equation 1:

$$
c_1 + \gamma V_1(\theta) = \mu(\theta)V_1'(\theta) + \frac{\sigma^2}{2}V_1''(\theta),
$$

with boundary conditions $V_1(\theta^*) = V_1'(\theta^*) = 0$. Now suppose μ is not constant, so strictly increasing somewhere. Construct a new pair of parameters $(\tilde{c}_1, \tilde{\mu})$ as follows: $\tilde{\mu}(\theta) = \mu(\theta_*)$ for all $\theta \in [\theta_*, \theta^*]$ and $\tilde{c}_1(\theta) = c_1(\theta) + (\tilde{\mu}(\theta) - \mu(\theta))V_1'(\theta)$ for all $\theta \in [\theta_*, \theta^*]$.[28] Then, by construction, the solution to Equation 1 is the same under these new parameters; in particular $\frac{V_1''(\theta_*)}{V_1'(\theta_*)}$. In addition, $\tilde{\mu}$ is non-decreasing (in fact constant), and $\tilde{c}_1$ satisfies $\tilde{c}_1(\theta) \geq \tilde{c}_1(\theta_*) = c_1(\theta_*)$ for all θ, with the inequality being strict at some θ. But then, by our previous discussion, if we instead take as our parameters $(\check{c}_1, \check{\mu})$ by $\check{c}_1 \equiv 1$ and $\check{\mu} = \tilde{\mu}$, we will attain a lower value of $\frac{V_1''(\theta_*)}{V_1'(\theta_*)}$.

The next step of the proof is to prove, under the assumption of constant c_1, c_2 and μ, that the tightest case is when $\mu = 0$.

Let $V(x)$ be defined as the solution to the ODE $1 + V(x) = \mu V'(x) + \frac{\sigma^2}{2}V''(x)$, with the initial conditions $V(0) = V'(0) = 0$. Note that this is the same as Equation (1) if we normalize $c \equiv 1$, $\gamma = 1$.

Disregarding the initial conditions, this ODE has a constant solution $V \equiv -1$, and the homogeneous ODE has general solution $k_1 e^{\alpha_1 x} + k_2 e^{\alpha_2 x}$, where α_1 and α_2 are the solutions

[28] The function V_1' used here is the solution to Equation 1 under the original parameters (c_1, μ).

to the quadratic equation $1 = \mu\alpha + \frac{\sigma^2}{2}\alpha^2$, i.e., $\alpha_1 = \frac{-\mu+\sqrt{\mu^2+2\sigma^2}}{\sigma^2}$, $\alpha_2 = \frac{-\mu-\sqrt{\mu^2+2\sigma^2}}{\sigma^2}$.

Since we require $V(0) = -1 + k_1 + k_2 = 0$ and $V'(0) = k_1\alpha_1 + k_2\alpha_2 = 0$, it follows that $k_1 = \frac{\alpha_2}{\alpha_2-\alpha_1}$ and $k_2 = -\frac{\alpha_1}{\alpha_2-\alpha_1}$. Then

$$V(x) = -1 + \frac{\alpha_2}{\alpha_2 - \alpha_1}e^{\alpha_1 x} - \frac{\alpha_1}{\alpha_2 - \alpha_1}e^{\alpha_2 x}.$$

Write $V_\mu(x)$ to make the dependence of V on μ explicit. Then what we aim to show is that, if $\mu > 0$, then $\frac{V'_\mu(x)}{V'_0(x)}$ is decreasing, and if $\mu < 0$, then $\frac{V'_\mu(x)}{V'_0(x)}$ is increasing.

Normalize $\sigma^2 = 2$. (We can do this by means of two changes of variables, rescaling time and the state space, respectively.) Then $\alpha_{1,2} = \frac{-\mu\pm\sqrt{\mu^2+4}}{2}$. For the sake of simplifying further, write $\tilde{\mu} = \frac{\mu}{2}$. Then $\alpha_{1,2} = -\tilde{\mu} \pm \sqrt{\tilde{\mu}^2 + 1}$. Finally, make the following change of variables: $y = e^x$, and denote $\hat{V} = V \circ \ln$. Then

$$V(x) = V(\ln(y)) = \hat{V}(y) = -1 + \frac{\alpha_2}{\alpha_2 - \alpha_1}y^{\alpha_1} - \frac{\alpha_1}{\alpha_2 - \alpha_1}y^{\alpha_2}.$$

We then aim to show that if $\mu > 0$, then $\frac{\hat{V}'_\mu(y)}{\hat{V}'_0(y)}$ is decreasing for $y \geq 1$, and that if $\mu' < 0$, then $\frac{\hat{V}'_{\mu'}(y)}{\hat{V}'_0(y)}$ is increasing for $y \geq 1$. We will use the following

Lemma 9. *Let $f, g : [a, b] \to \mathbb{R}$ be measurable positive functions such that $\frac{f(x)}{g(x)}$ is increasing for $x \in [a, b]$. Define $F, G : [a, b] \to \mathbb{R}$ as $F(x) = \int_a^x f(s)ds$, $G(x) = \int_a^x g(s)ds$. Then $\frac{F(x)}{G(x)}$ is increasing for $x \in [a, b]$.*

Proof. Write $\frac{f(x)}{g(x)} = \rho(x)$, and $\frac{F(x)}{G(x)} = \frac{\int_a^x g(s)\rho(s)ds}{\int_a^x g(s)ds}$. If $x' > x$, then

$$\frac{F(x')}{G(x')} = \frac{\int_a^{x'} g(s)\rho(s)ds}{\int_a^{x'} g(s)ds} = \frac{\int_a^x g(s)\rho(s)ds + \int_x^{x'} g(s)\rho(s)ds}{\int_a^x g(s)ds + \int_x^{x'} g(s)ds}.$$

Then $\frac{F(x')}{G(x')} \geq \frac{F(x)}{G(x)}$, since $\frac{\int_x^{x'} g(s)\rho(s)ds}{\int_x^{x'} g(s)ds} \geq \rho(x) \geq \frac{\int_a^x g(s)\rho(s)ds}{\int_a^x g(s)ds}$. $\square$

By this Lemma, it is enough to prove that if if $\mu > 0$, then $\frac{\hat{V}''_\mu(y)}{\hat{V}''_0(y)}$ is decreasing for $y \geq 1$, and that if $\mu' < 0$, then $\frac{\hat{V}''_{\mu'}(y)}{\hat{V}''_0(y)}$ is increasing for $y \geq 1$. Now note that

$$\hat{V}_0(y) = -1 + \frac{y}{2} + \frac{1}{2y}, \quad \hat{V}''_0(y) = \frac{1}{y^3}$$
$$\hat{V}'_\mu(y) = \frac{\alpha_1\alpha_2}{\alpha_2 - \alpha_1} y^{\alpha_1 - 1} - \frac{\alpha_1\alpha_2}{\alpha_2 - \alpha_1} y^{\alpha_2 - 1} \propto y^{\alpha_1 - 1} - y^{\alpha_2 - 1}$$
$$\hat{V}''_\mu(y) \propto (\alpha_1 - 1) y^{\alpha_1 - 2} - (\alpha_2 - 1) y^{\alpha_2 - 2}$$
$$\frac{\hat{V}''_\mu(y)}{\hat{V}''_0(y)} \propto (\alpha_1 - 1) y^{\alpha_1 + 1} - (\alpha_2 - 1) y^{\alpha_2 + 1}$$

We can verify that, if $\mu > 0$, $\alpha_1 \in (0,1)$ and $\alpha_2 < -1$, so $\alpha_1 - 1 < 0$, $\alpha_1 + 1 > 0$, $\alpha_2 - 1 < 0$ and $\alpha_2 + 1 < 0$. Hence $\frac{\hat{V}''_\mu(y)}{\hat{V}''_0(y)}$ is decreasing in y. Similarly, if $\mu < 0$, $\alpha_1 > 1$ and $\alpha_2 \in (-1, 0)$, so $\alpha_1 - 1 > 0$, $\alpha_1 + 1 > 0$, $\alpha_2 - 1 < 0$ and $\alpha_2 + 1 > 0$. Hence $\frac{\hat{V}''_\mu(y)}{\hat{V}''_0(y)}$ is increasing in y. This finishes the proof.

Finally, note that in the case where $c_1 \equiv c_2 \equiv 1$ and $\mu \equiv 0$, the result is (weakly) trivially true, as $H_1(\theta_*, \theta^*) \equiv H_2(\theta_*, \theta^*)$ so that $\frac{|\frac{\partial H_1(\theta_*, \theta^*)}{\partial \theta_*}|}{|\frac{\partial H_2(\theta_*, \theta^*)}{\partial \theta_*}|} \equiv 1$. It follows from our arguments that, if c_1 and c_2 are strictly increasing (decreasing), then $\frac{|\frac{\partial H_1(\theta_*, \theta^*)}{\partial \theta_*}|}{|\frac{\partial H_2(\theta_*, \theta^*)}{\partial \theta_*}|}$ will be strictly increasing instead. □

Proof of Proposition 4. Parts (i) and (ii) can be proved in the same fashion as their analogues in Proposition 2.

Part (iii) is a consequence of Lemma 8. Indeed, suppose WLOG that $\mu \leq 0$, and H_1, H_2 are increased to H'_1, H'_2 with $\frac{H'_1}{H_2} = \frac{H_1}{H_2}$. We want to show that θ^* increases. This is equivalent to showing that, if H_2 is increased to H'_2, the unique value $\tilde{H}_1 > H_1$ that leaves θ^* unchanged satisfies $\tilde{H}_1 < H'_1$, i.e., $\frac{\tilde{H}_1}{H_2} < \frac{H_1}{H_2}$.

This is equivalent to showing that $\frac{H_1(\theta_*, \theta^*)}{H_2(\theta_*, \theta^*)}$ is increasing in θ_*. But that follows from the fact that $\frac{\frac{\partial H_1(\theta_*, \theta^*)}{\partial \theta_*}}{\frac{\partial H_2(\theta_*, \theta^*)}{\partial \theta_*}}$ is increasing in θ_* (Lemma 8) combined with Lemma 9. □

Proof of Proposition 8. We provide only a sketch of the proof, as it is largely analogous to the proof of Proposition 1.

The first step is to note that the game is supermodular, in the following sense: order the players' strategy spaces so that a "higher" strategy for player 1 is one with a higher surrender probability at every history, and a "higher" strategy for player 2 is one with a lower surrender probability at every history. Then, given two strategies $\psi \geq \psi'$ for player 1, player 2's best response to ψ must be weakly higher than her best response to ψ'. The proof is similar to that of Lemma 1. Briefly, the reason is that player 2's equilibrium payoff at any history must be weakly higher when facing a player 1 more likely to surrender, and up to indifference, player 2 ought to surrender iff her payoff from continuing is negative. Then all equilibria are bounded in this sense between a greatest and a smallest equilibrium.

Call a strategy for player i (anti)monotonic if it has player i surrender whenever θ falls in a (anti)monotonic set A. We refer to such a strategy simply by its surrender set A. The second step is to note that, if player 1 plays a monotonic strategy A, then player 2's best response is an antimonotonic surrender set $BR_2(A)$; conversely, if player 2 plays an antimonotonic set B, player 1's best response $BR_1(B)$ is monotonic. It follows that, in both the greatest and the smallest equilibrium of the game, player 1's strategy is monotonic and player 2's is anti-monotonic.

The third step is to show that the greatest and the smallest equilibrium are identical (up to measure zero). Denote player 1's and 2's surrender regions in the greatest equilibrium by A, B respectively, and their surrender regions in the smallest equilibrium by A', B', respectively. By assumption, $A \supseteq A'$ and $B \subseteq B'$. Assume that at least one of these inclusions is strict (otherwise we are done).

It can be shown that the players' best-response mappings are contractions in a certain sense. Namely, given two monotonic sets $C \supseteq C'$, let $d(C, C') = \inf\{a : C + av \subseteq C'\}$.[29] Define d analogously for nested pairs of antimonotonic sets. Then, using Assumptions A2' and B1', it can be shown that if $d(C, C') > 0$ then $d(BR_2(C), BR_2(C')) < d(C, C')$ for $C \supseteq C'$ monotonic, and analogously for antimonotonic sets. The key point is as follows: suppose $\theta \in \partial BR_2(C)$ for some monotonic set C (i.e., player 2 is indifferent about surrendering in state θ when C is player 1's surrender region). Then she must strictly prefer to continue in state $\theta + av$ when player 1's surrender set is $C + av$, for any $a > 0$, due to assumptions A2' and B1'. It follows that $BR_2(C + av) \subsetneq BR_2(C) + av$ for any $a > 0$. In particular, $BR_2(C') \subseteq BR_2(C + d(C, C')v) \subsetneq BR_2(C) + d(C, C')v$ for any $C \supseteq C'$ monotonic. Moreover, the closure of $BR_2(C + d(C, C')v)$ must be contained in the interior of $BR_2(C) + d(C, C')v$. From here, it follows[30] that $d(BR_2(C + d(C, C')v, BR_2(C) + d(C, C')v) > 0$ and hence $d(BR_2(C), BR_2(C')) < d(C, C')$. The argument is analogous for player 1.

But then $d(B, B') < d(A, A') < d(B, B')$, a contradiction.[31] That A and B are disjoint follows from Assumption B6'.

□

[29] Given a set S and a vector v, we denote $S + v = \{s + v : s \in S\}$.
[30] Here, we use that $\prod_{i=1}^{k}[-M_i, M_i]$ is compact.
[31] Assumption B3' guarantees that $d(A, A')$, $d(B, B')$ are finite.

C Two-Sided Concessions (For Online Publication)

In this Section we briefly discuss the case in which both players are able to make concessions, and can do so at any time. (As we will see, in this case, timing will matter.) This setting can be taken as modeling bargaining under an extreme lack of commitment, so that, while both players can give up things, it is impossible for them to make *quid pro quo* bargains to resolve matters.

To simplify the analysis, we assume that the prize is only finitely divisible, that is, there is a finite sequence $0 = x_0 < x_1 < \ldots < x_k = 1$ such that each interval $[x_l, x_{l+1})$ cannot be split. Thus, at any given time, the part of the prize still in dispute must be of the form $[x_l, x_j)$ ($l < j$), if player 1 has conceded $[x_j, 1]$ and player 2 has conceded $[0, x_l)$. To ensure that it is never optimal to concede non-consecutive intervals, we will assume that $v_1^1 > \ldots > v_1^k$ and $v_2^1 < \ldots < v_2^k$, where $v_i^l = \int_{x_{l-1}}^{x_l} v_i(x)dx$.

The relevant state of the game is now of the form (θ, x_l, x_j). For simplicity, we will focus on threshold strategies. A threshold strategy for player 1 is given by a collection of thresholds $\theta^*(x_l, x_j)$ for each $l < j$ such that, if the prize still in dispute is $[x_l, x_j)$, then player 1 concedes the next part of the prize—thus changing the state to (θ, x_l, x_{j-1})—as soon as $\theta_t \geq \theta^*(x_l, x_j)$. The definition for player 2 is analogous.[32] Denote by $\tilde{\theta}_*(x_l, x_j)$, $\tilde{\theta}^*(x_l, x_j)$ the equilibrium surrender thresholds if the prize in dispute is $[x_l, x_j)$ and *no (further) partial concessions are allowed*. (These "naive" thresholds are unique and pinned down by our analysis in Section 2.)

An equilibrium in threshold strategies can be constructed as follows. When only one part of the prize is in dispute (i.e., $j - l = 1$), the game is equivalent to the baseline model, for which the solution is known. When two parts are in dispute ($j - l = 2$), the players are still effectively facing a war of attrition, but with more complicated payoffs after (partial) victory or surrender: at the moment when player 1 makes a concession in state θ, the players' continuation payoffs are $(V_1(\theta, x_l, x_{j-1}), V_2(\theta, x_l, x_{j-1}) + v_2^j)$ rather than $(0, H_2)$, and analogously for player 2. Finding equilibrium thresholds for this game pins down value functions $V_i(\theta, x_l, x_j)$ for $j - l = 2$. We can then consider states with $j - l = 3$, and so on.[33]

Giving a general and explicit equilibrium characterization is difficult. A major reason is

[32] We restrict the players to conceding one part of the prize at a time for simplicity. It makes little difference, as multiple concessions can be made at once—in particular if $\theta^*(x_l, x_{j-1}) \leq \theta^*(x_l, x_j)$ or $\theta_*(x_l, x_{j-1}) \geq \theta_*(x_l, x_j)$. Also, we assume that if both players try to concede at the same time, each one succeeds (first) with probability 0.5. Again, this makes little difference.

[33] It can be shown at each step that (a) best response to a threshold strategy is another threshold strategy; the mapping giving the optimal threshold(s) for each player as a function of the opponent's concession threshold is a upper-hemicontinuous and convex-valued correspondence; and the composition of the two must cross the identity, so an equilibrium in threshold strategies exists.

that, with two-sided concessions, the equilibrium concession(s) are no longer predetermined as a function of the fundamentals, as in Proposition 6; instead, they may depend on the entire realized path of θ_t.

We show this in the special case $k = 2$, i.e., when the prize is made up of only two pieces. To simplify notation, denote the disputed region in the model without concessions by $[\tilde{\theta}_*(0,1), \tilde{\theta}^*(0,1)] = [\underset{\sim}{\theta}, \tilde{\theta}]$; the disputed region after a concesion by player 1 by $[\tilde{\theta}_*(0,x_1), \tilde{\theta}^*(0,x_1)] = [\theta_1, \theta^1]$; the disputed region after a concession by player 2 by $[\tilde{\theta}_*(x_1,1), \tilde{\theta}^*(x_1,1) = [\theta_2, \theta^2]$; and the equilibrium concession thresholds by $\theta_*(0,1) = \theta_*$, $\theta^*(0,1) = \theta^*$. Assume parameters such that $\theta^1 > \theta^0$ and $\theta_2 < \theta_0$, i.e., making a concession would be wortwhile for either player in the one-sided version of the game (Proposition 6).

Figure 4: Concessions under highly heterogeneous valuations

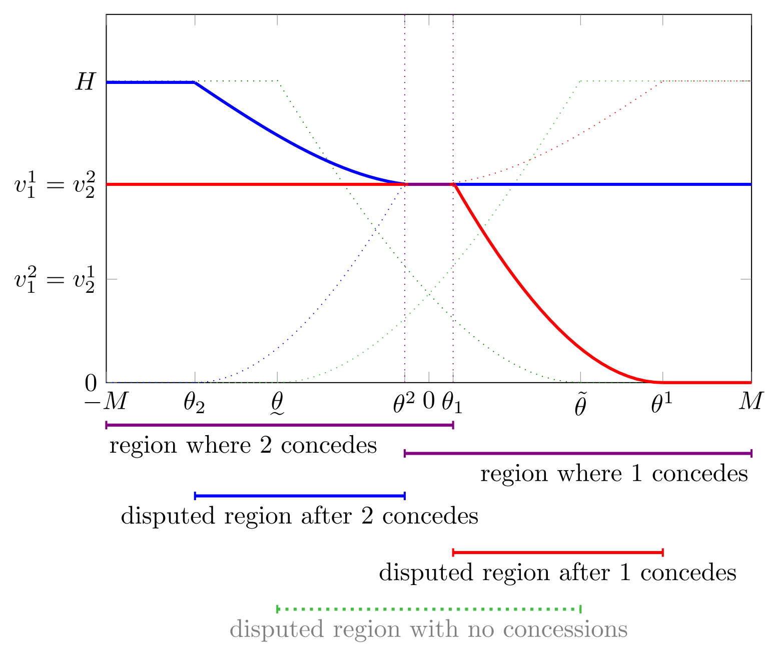

Suppose first that $\theta^2 \leq \theta_1$. (Intuitively, this condition holds with the players' valuations of the two parts of the prize are different enough.) Then, for $\theta_t \in [\theta^2, \theta_1]$, the game ends immediately, with player 1 getting $[0, x_1)$ and 2 getting $[x_1, 1]$. Indeed, this is the outcome if either player makes a partial concession, and it must generate higher payoffs than fighting for at least one player. Then, when $\theta_t \in (\theta_1, \theta^1)$, the players effectively fight over player 1's "turf", $[0, x_1)$, with $[x_1, 1]$ guaranteed to go to player 2; conversely, when $\theta_t \in (\theta_2, \theta^2)$, the players fight only over player 2's turf, $[x_1, 1]$. The game thus devolves into one of two possible wars over individual pieces of the prize. Figure 4 illustrates an example with $c_1(\theta) = 5 + \frac{10}{3}\theta$ $c_2(\theta) = 5 - \frac{10}{3}\theta$, $\sigma^2 = 2$, $v_i^i = 7$, $v_i^j = 3.5$ and $H_i = 10.5$. The red curve is player 1's continuation value after a concession, $V_1(\theta, 0, x_1)$, and the blue curve is her continuation value if player 2 concedes to her, $V_1(\theta, x_1, 1) + v_1^2$.

Figure 5: Concessions under moderately heterogeneous valuations

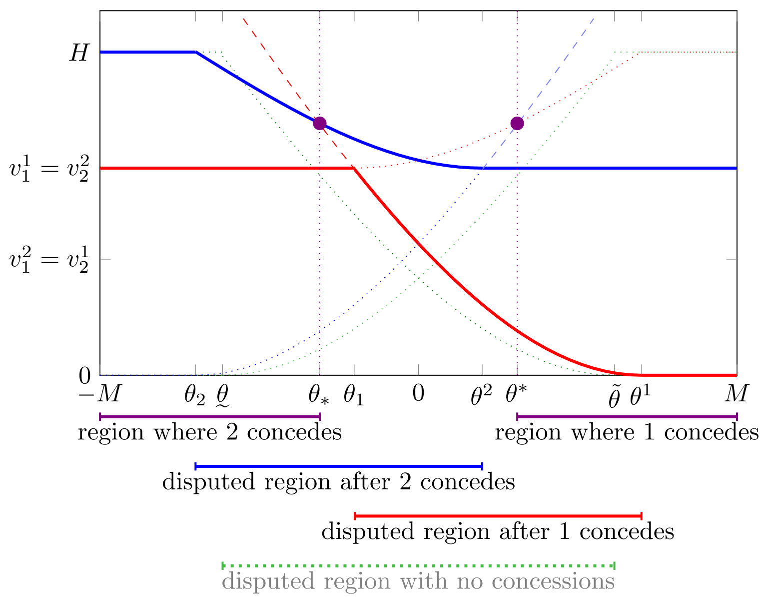

The more interesting case is when $\theta^2 > \theta_1$, i.e., valuations are only moderately heterogeneous, as illustrated in Figure 5 under the parameter values $v_i^i = 12.5$, $v_i^j = 7$, $H_i = 19.5$ and the same cost functions as in the previous figure.[34] In equilibrium, player 1 makes a concession when θ_t first goes above a threshold value θ^*, while player 2 makes a concession if θ_t goes below θ_*.

Since θ^* lies in the interior of $[\theta_1, \theta^1]$, a concession by 1 at θ^* does not lead to surrender by either player, and likewise for player 2. As a result, the outcome of the war depends on the entire path followed by the state. For instance, suppose $\theta_0 = 0$. Then, if θ_t first goes below θ_2 before ever turning positive, player 2 gives up everything—first with a concession at $\theta_t = \theta_*$ and then with surrender at θ_2. However, if θ_t first goes up beyond θ^* (but below θ^1) and then down towards θ_2, player 1 concedes $[x_1, 1]$ when the state crosses θ^*, and then player 2 ends the war by giving up $[0, x_1]$ when θ_t goes below θ_1.

The threshold θ_* is chosen to make player 1 indifferent about conceding for all $\theta \geq \theta_1$; similarly, θ^* leaves 2 indifferent about conceding for $\theta \leq \theta^2$. Other threshold equilibria are not possible. Indeed, we can show that, were player 2's concession threshold any lower, player 1 would strictly prefer to concede at all $\theta \geq \theta_1$ (i.e., $\theta^* \leq \theta_1$), whence a marginal concession by 1 would lead to surrender by 2; thus, between θ_* and θ^*, the players would be

[34] Of course, if valuations are too homogeneous, then concessions are weak by Proposition 4.(iii). It is not hard to show that there is a nonempty range of prize valuations with intermediate heterogeneity for which concessions are not weak but the intervals (θ_2, θ^2), (θ_1, θ^1) overlap.

effectively fighting over $[x_1, 1]$, as $[0, x_1)$ is guaranteed to go to player 1. But then we ought to have $[\theta_*, \theta^*] = [\theta_2, \theta^2]$, contradicting $\theta^* \leq \theta_1$. Similarly, if θ_* were any higher, player 1 would not concede for any $\theta \leq \theta^1$, whence player 2 would concede for all $\theta \leq \theta^2$, leading to the same contradiction.