Colonel Blotto's Tug of War*

Tilman Klumpp†

June 2021

Abstract

We examine Tug of War contests with the Blotto specification. Players have fixed effort budgets and must allocate these budgets to a sequence of battles. The outcome of each battle is a function of the efforts allocated to that battle. The player who first wins L more battles than the opponent wins the contest. We prove the one-step deviation principle for the undiscounted version of this game. We then derive a pure strategy, subgame perfect equilibrium for the case where the contest success function that governs each battle is a generalized Tullock function with exponent $1/2$ or less. In the equilibrium, the players invest the same percentage of their remaining resources into each battle. The value of this percentage depends on how close each player is to winning the contest. Escalation of efforts, measured in relation to the players' remaining budgets, occurs when the player with the smaller budget is close to winning. At the same time, the probability that a player wins any individual battle remains constant along the entire equilibrium path.

Keywords: Tug of War games; Colonel Blotto games; sequential contests; multi-battle contests.

JEL codes: D72; D74.

---

*I thank Ashantha Ranasinghe and Xuejuan Su for helpful comments, and especially Kai Konrad for suggesting this topic. Financial support from the Social Sciences and Humanities Research Council of Canada (grant IG-435-2015-1397) is greatly acknowledged.

†Department of Economics, University of Alberta. 8-14 Tory Building, Edmonton, AB, Canada T6G 2H4. E-mail: klumpp@ualberta.ca.

1 Introduction

Many strategic situations in business, politics, sports, and other fields can be regarded as multi-battle contests. These are games in which players compete by investing resources into a series of battles, the outcome of each battle depends on these investments, and payoffs depend on how many and/or which battles each player wins. This competition is often dynamic, in which case players can condition their investments in one battle on the outcomes of previous battles.

Table 1 below categorizes two-player, sequential, multi-battle contests along two dimensions.[1] The first dimension focuses on the players' objective: In a majoritarian contest the first player to win L battles wins, whereas in a *Tug of War* the first player to win L more battles than the opponent wins. The second dimension focuses on the players' resources: In adjustable budget (or cost-of-effort) models, players can invest any amount of effort; however, effort is costly and the total cost of effort is deducted from a player's utility. In fixed budget models, effort has no such direct cost; instead, players are given finite effort budgets that they cannot exceed. This latter specification is also called a *(Colonel) Blotto* model. As Table 1 shows, the literature on sequential multi-battle contests remains incomplete insofar as Tug of War games with the Blotto specification have not been studied. The present paper aims to fill this gap.

Table 1: Models of sequential multi-battle contests.

	Majoritarian objective *("First to win L battles")*	**Tug of War objective** *("First to win L more battles")*
Adjustable budgets *(Cost-of-effort model)*	Klumpp/Polborn (2006) Konrad/Kovenock (2009) Malueg/Yates (2010) Sela (2011) Gelder (2014) Fu/Lu/Pan (2015)	Harris/Vickers (1987) Budd/Harris/Vickers (1993) Konrad/Kovenock (2005) Agastya/McAfee (2006) Häfner (2017, 2020)
Fixed budgets *(Blotto model)*	Klumpp/Konrad/ Solomon (2019)	**This paper**

We examine a Tug of War game in which each contestant is endowed with an initial effort budget. Players simultaneously decide how much of their budget to invest in the first battle. A winner of the first battle is then drawn, with the win probabilities depending on

[1] Fu, Lu, and Pan (2015), Häfner (2017), and Häfner (2020) study team contests played between two groups, each consisting of many decision making agents. A detailed review of the papers listed in Table 1 will be given in Section 2.

the players' efforts. The players observe each others' efforts and the outcome of the first battle before deciding how much of their remaining budgets to invest in the second battle. This process continues until one player has accumulated $L \geq 2$ more battle victories than the opponent for the first time. At this point, the player who is ahead wins the contest. Any unspent resources have no residual value.

As an application that can be described by our model, consider two startups companies that have each received funding from investors and are now entering a large number of local markets. In each local market, the firms provide a regulated service, such as electric scooter rentals or a ride sharing app. Firms enter these markets in a common sequence, determined by the order in which local governments permit these services to operate. Each local market is small enough for only one firm to be viable, and an entrant's success in a given market depends on how much both competitors spend on staffing and advertising in that market. Investors of a firm whose market share falls behind that of its rival by a large enough margin prefer to sell the venture to the rival (instead of providing additional funding to continue competition). On the other hand, the decisions of how many resources to allocate to each local market are made by the startups' founders, who only want their respective firms' to win. Thus, the founders derive no value from preserving funds in either outcome of the race.

The literal application of our model is, of course, the actual Tug of War game—a strength contest that was an Olympic event until 1920 and is still actively played in many parts of the world. The Blotto budget specification is realistic if athletes have limited physiological resources and and if competitions are sufficiently far apart temporally for players to not gain an advantage from preserving resources in any given match.[2] Finally, as their name suggests, Colonel Blotto games are often motivated by military applications in which two sides must allocate limited resources across multiple battlefields. The Tug of War objective, then, represents conflicts in which battles are fought sequentially and ultimate victory depends on winning sufficiently many more battles than the adversary. The Blotto assumption that resources have value outside of the model is often realistic in these applications. For example, there are no civilian uses for mortar shells, but a mortar shell used in one military battle cannot be used in another battle (generating an opportunity cost within the model).

We identify a pure strategy, subgame perfect equilibrium of the Blotto Tug of War for the cases where the winner of each battle is determined through a generalized Tullock contest success function (Tullock 1980). In the equilibrium, both players invest a common percentage of their remaining resources into each battle, leaving them with the same

[2]Other sporting contests, too, incorporate elements of the Tug of War. Under the advantage scoring rule in tennis, for example, if both players have scored three points ("40 all") the game becomes a Tug of War with $L = 2$.

relative budgets in the next battle. This implies that the probability that a player wins any given battle stays constant throughout the Tug of War, in equilibrium—a property the Blotto Tug of War shares with the Blotto majoritarian contest (see Klumpp, Konrad, and Solomon 2019). By the same token, the lead a player has over his opponent follows a simple random walk—a property the Blotto Tug of War shares with the some non-Blotto Tug of War contests (see Häfner 2017). We show that the fraction invested by both players is equal to the inverse of the expected remaining duration of the contest in the *opposite* state of the game. For example, suppose player A is two victories away from winning the Tug of War. In this case, both players divide their remaining budgets by the expected remaining duration of the same contest, but assuming that player B was two victories away from winning, and invest the resulting ratio in the current battle. If the players start with identical budgets, this strategy implies that players escalate their efforts the higher one player's lead over their opponent—another property the Blotto Tug of War shares with some non-Blotto Tug of War models (see, e.g., Agastya and McAfee 2006).[3]

Despite these similarities, the Blotto Tug of War contest presents analytical challenges not present in either Tug of War contests with adjustable budgets or majoritarian contests with fixed budgets. Most importantly, among the contest formats in Table 1, the Blotto Tug of War is the only contest in which players must allocate a finite effort budget to potentially infinitely many battles. This implies that certain "simple" strategies, such as spreading one's budget equally across all battles (which would be the equilibrium of the majoritarian Blotto contest), are unavailable. More generally, strategies cannot prescribe the same positive effort at all battles in which a player is the same number of previous victories ahead of his opponent (which would be the case in Markovian equilibria of the non-Blotto Tug of War contest). To obtain useful results in an infinitely long game, however, some degree of self-similarity across subgames is clearly necessary. To this end, certain assumptions on the contest technology and on payoffs must be made. In particular, the assumption of a Tullock contest success function, which is homogeneous of degree zero, ensures that any two subgames in which (i) player A has the same lead over player B and (ii) player A's remaining budget *relative* to player B's remaining budget is the same, are scaled versions of one another. Thus, the continuation equilibrium in one subgame is a scaled version of equilibrium play in the other.

In order to fully translate this fractal structure into a solution, two more restrictions are necessary still. First, we assume that players do not discount future payoffs. This

[3] If players start with asymmetric budgets, percentage efforts escalate sharply when the poorer player establishes a lead (as the expected remaining duration in the scenario where the richer player had the lead would decrease); if the stronger player establishes a lead, percentage efforts first decrease, but then increase again when the lead becomes sufficiently large.

assumption ensures that continuation payoffs can always be expressed in terms of the gambler's ruin problem, which facilitates a closed-form solution for the equilibrium. The disadvantage is that the undiscounted Tug of War contest is not continuous at infinity, which would be the "usual" sufficient condition to ensure that the one-step deviation principle holds. Therefore, we have to establish the one-step deviation principle for the Blotto Tug of War from the ground up (without relying on continuity at infinity via discounting). Second, while we obtain the equilibrium pure strategy profile analytically based on first-order conditions, the requisite higher-order conditions can only be verified numerically. This verification shows the first-order conditions to be sufficient only when the exponent in the Tullock function is 1/2 or less. In other words, the probability of success in a battle must not only be concave in own effort, but must be sufficiently concave, for our pure strategy equilibrium to exist. The concavity requirement we identify is stricter than in previously examined contests.[4]

The remainder of the paper is organized as follows. After reviewing the related literature in Section 2, we describe our formal model of the Blotto Tug of War in Section 3. In Section 4, we prove the one-step deviation principle for the undiscounted version of this game. In Section 5, we derive a candidate strategy profile for subgame perfect equilibrium analytically, and verify the higher-order equilibrium conditions numerically. Section 6 discusses the dynamics of players' efforts in equilibrium, and Section 7 concludes.

2 Related Literature

Multi-battle contests with fixed budgets have a long history in economics and statistics. Static Blotto models were first analyzed in Borel (1921), and more recently in Roberson (2006), Kvasov (2007), and Roberson and Kvasov (2012). In these contests, players allocate effort to all battlefields at the same time, without being able to condition effort in one battle on the outcomes of previous battles. Since this paper examines a dynamic contest, where such conditioning is possible and an essential part of equilibrium play, we will not review the literature on static Blotto games and focus on sequential contests only.[5]

Table 1 lists contributions to the literature on sequential multi-battle contests, focusing on papers that examine models with either the majoritarian or Tug of War objective.

[4] For example, in the one-shot Tullock contest with constant marginal cost of effort, a pure strategy equilibrium exists if the exponent on effort is 2 or less, which includes even some non-concave functions (see Baye, Kovenock, and De Vries 1994). In the Blotto majoritarian contest examined in Klumpp, Konrad, and Solomon (2019), a pure strategy equilibrium exists if the contest success function is a Tullock function with exponent 1 or less.

[5] A general introduction to the theory of static and dynamic contests is Konrad (2009). Kovenock and Roberson (2010) review specifically the literature on multi-battle contests, including static Blotto contests.

Most of these models differ from ours in that they are of the adjustable-budget type. That is, players are not constrained in how much they can invest, but must pay some utility cost for the effort they choose to invest. This specification captures situations in which effort not invested in the contest has alternative uses *outside of* the model, with the cost of effort reflecting the value of effort in the next best (unmodeled) alternative use. In contrast, under the Blotto specification there are no such alternative uses outside of the model, leaving the opportunity costs that arise implicitly *within* the Blotto contest as the only cost of effort.

2.1 Tug of War contests

The first game theoretic models of the Tug of War is Harris and Vickers (1987). There, two players compete by investing efforts in consecutive battles. The outcome of each battle is determined through a stochastic process closely related the Tullock contest success function. Each player pays a cost of effort that depends on his effort level and the time until the battle is resolved. In equilibrium, the leader spends more than the follower and is more likely to win subsequent battles. At the same time, the effort by the player who has fallen behind decreases in the opponent's lead.

Budd, Harris, and Vickers (1993) develop a continuous time duopoly model in which firms affect the evolution of their market shares by investing costly resources. Firms earn product market profits that depend on their current market share, and may also receive additional benefits once a boundary state is reached in which the rival's market share drops to zero. The Tug of War is a special case of this model, where firms earn zero product market profits until this boundary event occurs. Firms maximize the present value of their product market profits, minus the cost of resources they invest in building market share. Budd, Harris, and Vickers (1993) identify a number of forces that govern the firms' investments in equilibrium. Most importantly, the leader invests more than the laggard, and is able to extend its lead, if the firms' *joint* profits increase in the lead—which is the case in the Tug of War formulation of their model.

Konrad and Kovenock (2005) and Agastya and McAfee (2006) study Tug of War games in which each battle is an all-pay auction and effort has a constant per-unit payoff cost. The two models differ in a number of aspects and result in different equilibrium behavior. Konrad and Kovenock (2005) assume that battles in which both players exert the same effort are resolved in favor of the player who has the higher continuation payoff under the assumption that he wins all subsequent battles at zero effort. In the Markovian equilibrium, players invest effort only at one or two states of the game (a state is characterized by the lead of one player over the other). Agastya and McAfee (2006), on the other hand, assume that battles in which both players exert the same effort are resolved randomly, and that battles in which both player exert zero effort must be refought. In addition, they assume a negative prize for losing the Tug of War. Because there is discounting, the player who has fallen behind is primarily motivated by avoiding or delaying loss, whereas the player in the lead is motivated by speeding up victory. Depending on the prize for winning relative to the penalty for losing, two types of equilibrium emerge. In one, efforts escalate as a player's lead over the other widens. In the other, fighting eventually ceases forever without a winner or loser being determined (an "eternal peace" equilibrium). As discussed previously, we assume no discounting in our Blotto model in order to obtain a closed form solution. At the same time, we assume a penalty of losing, but because there is no discounting, this does not translate into an incentives to delay a loss for any finite number of periods.[6]

2.2 Majoritarian contests

Closely related to Tug of War contests are sequential majoritarian contests. In such games, there are $2L - 1$ potential battles, and the player who first wins L battles wins the contest. Klumpp and Polborn (2006) study a majoritarian multi-battle contest with constant marginal cost of effort, assuming each battle is a Tullock contest.[7] In pure strategy equilibrium of this contest, early wins create "momentum:" For a player who has fallen behind, the incentive to invest additional resources in order to catch up is weakened, making it relatively easy for the frontrunner to increase his lead. At the same time, the total effort invested by both players decreases in the frontrunner's lead. Konrad and Kovenock (2009) examine a similar sequential majoritarian structure, but assume that each battle is an all-pay auction.[8] In this case a pure strategy equilibrium no longer exists. In the unique mixed strategy equilibrium, players invest resources only at a small number of states of the game (unless intermediate prizes for individual battle victories are awarded). Generally, effort is non-monotonic in the frontrunner's lead. Gelder (2014) extends this model by introducing both discounting and a penalty for losing. Similar to the Tug of War model in Agasty and McAfee (2006), the players who falls behind spends effort to delay the loss, resulting "last stand" behavior.

Klumpp, Konrad, and Solomon (2019) examine a sequential contest with the majoritarian objective and the Blotto specification, assuming a smooth contest success function (but not necessarily the Tullock function). The unique subgame perfect equilibrium is, at

[6]The losing penalty does make players prefer a perpetual Tug of War to one that ends in a loss; however, in the equilibrium, this outcome is a null event and all results would continue to hold if the losing penalty was zero.

[7]Malueg and Yates (2010) examine a similar model, restricted to $L = 2$, and empirically test it using data from tennis matches.

[8]Sela (2011) examines a similar sequential all-pay auction model, restricted to $L = 2$.

each stage, for both players to divide their remaining budgets evenly across the remaining battles. Thus, the success probabilities do not change from battle to battle, and efforts do not change once a player establishes a lead over the opponent. The equilibrium we find in the Blotto Tug of War shares the first feature. However, as discussed above, players cannot achieve this by dividing their budgets "evenly" across battles. As a result, the actual effort allocation strategies are considerably more complex than in the majoritarian contest, and involve effort adjustments both over time and in response to one player establishing a lead over the opponent. For the case of symmetric budgets, the equilibrium efforts can be interpreted as players stretching their remaining budgets across the *expected* remaining duration of the contest.[9]

2.3 Sequential multi-battle team contests

Finally, the recent literature has examined team versions of both the majoritarian and Tug of War non-Blotto contests. Team contests are multi-battle contests fought between two groups, each consisting of several members. Every battle is contested by a pair of individuals from each groups. While the final prize received by the winning group is a public good among its members, the efforts are privately costly to the individuals.

Fu, Lu, and Pan (2015) examine a sequential majoritarian team contest for a wide range of success functions. In equilibrium, the efforts invested in any given battle depend on the characteristics of the players fighting this battle, and on the characteristics of the battle (e.g, the values of intermediate prizes, if any), but *not* on the state of the overall multi-battle contest (i.e., on how many previous victories the teams have accumulated). If all members of a team have the same cost functions and valuations, the probability that a given team wins each battle is a constant. This finding is in contrast to the results in Klumpp and Polborn (2006) and Konrad and Kovenock (2009), but mirrors the result in Klumpp, Konrad, and Solomon (2019).

Häfner (2017) examines a Tug of War team contest, assuming that each battle is an all-pay auction. Häfner (2017) allows individual team members to differ in their valuations of their team winning the Tug of War, but assumes that these valuations are private knowledge. In the equilibrium, while efforts depend on the state, the win probabilities in each battle do not. Thus, the probability that a team wins any given battle remains constant, regardless of how close the team is to winning or losing the Tug

[9] In general, the equilibrium strategy calls for players to stretch their remaining budgets across the expected remaining duration of the contest in the "opposite" state (i.e., if player A's lead was instead player B's lead). This duration is the same as the actual remaining duration if and only if the two players' remaining budgets are the same.

of War. This finding is in contrast to the results in Konrad and Kovenock (2005) and Agastya and McAfee (2006).[10] However, it is similar to the result in this paper.

Thus, sequential team contests with adjustable budgets and the all-pay auction success function, and sequential individual-player contests with Blotto budgets and the Tullock success function, give rise to similar equilibrium dynamics. In particular, in both cases, win probabilities in individual battles stay fixed throughout the duration of the contest.

3 Tug of War Games with Fixed Resource Budgets

Consider the game of *Tug of War* between players A and B. Time is discrete with $t = 1, 2, \ldots$. In each period $t \geq 1$ players A and B compete in a battle, which is either won or lost, until the time at which a player has won L more battles than his opponent for the first time, where $L > 1$. In this event, the game ends and the player who is up L victories over his opponent wins a prize of $+1$, while the opponent obtains -1. There is no discounting, and prior to the game ending each player obtains a flow payoff of 0 per period.

Each battle is governed by a contest success function $f : \mathbb{R}^2_+ \to [0, 1]$, where $f(x, y)$ is the probability that A wins the battle if A spends effort x and B spends effort y, on the given battle. The probability that B wins the battle is $1 - f(x, y)$. In Section 5 we will assume a specific parametric form for f. However, at present we only need to assume the following minimal properties: f is increasing in its first argument and decreasing in its second argument, $f(0, y) = 1 - f(x, 0) = 0$ for all $x, y > 0$, and $0 < f(x, y) < 1$ for all $x, y > 0$.

We assume that players are endowed with initial effort budgets $\bar{a} > 0$ and $\bar{b} > 0$ which they can invest in the battles. A player cannot spend more than his initial budget in total, and in each battle he cannot use more than the difference between his initial budget and the effort already spent. Players observe the outcome of each battle and the opponent's remaining resources before making (simultaneous) investment decisions into the next battle. Unspent effort budgets have no value once the game is over.

3.1 State space

It will be convenient to introduce the notion of a state that summarizes the outcomes of past battles.

Entering period $t \geq 0$, let $0 \leq S_t^A \leq t$ be the number of previous battles that A has won, and let $0 \leq S_t^B \leq t$ be the number of previous battles that B has won. We define

[10] In a related paper, Häfner (2020) examines a team version of Tug of War contest that is otherwise the same as the contest examined in Agastya and McAfee (2006), and shows that an eternal peace equilibrium no longer exists.

a *state* of the tournament as a tuple $(t,d) \in \mathbb{N} \times \{-L, \dots, L\}$, where $d = S_t^A - s_t^B$. If $d \in \{-L, L\}$, then (t,d) is a *terminal state.* Once a terminal state (t,d) is reached, A wins the prize if $d = L$ and B wins the prize if $d = -L$. If $-L < d < L$, then (t,d) is a *battle state.* At each battle state (t,d), if A wins the battle then the next state is $(t+1, d+1)$, and if B wins the battle the next state is $(t+1, d-1)$.

Figure 1 illustrates the state space for the the case $L = 3$: The states are arranged in a lattice. Player A winning a battle corresponds to a step up and to the right in this lattice, and player B winning a battle corresponds to a step down and to the right. A sample path is shown in which A wins the first battle, loses the second and third, and then wins the four consecutive battles (and thus the game).

Figure 1: State space and a sample path in Tug of War with $L = 3$.

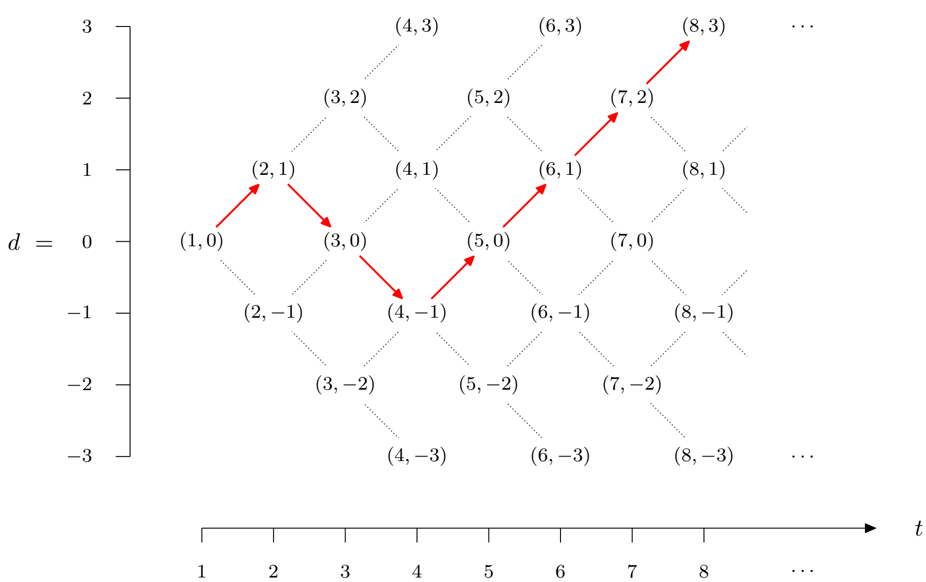

Given battle state (t,d), let $S(t,d) \equiv \{(t', d') : t' > t,\ -L < d' < L,\ |d' - d| \leq t' - t\}$ denote the set of battle states that can be reached from (t,d). Set $W(t,d) \equiv S(t,d) \cup (t,d)$ and note that the set of all battle states of the Tug of War is $W(1,0)$.

3.2 Strategies and equilibrium

An *investment function* for player A at battle state (t,d) is a mapping

$$\alpha_{t,d} : [0, \bar{a}] \times [0, \bar{b}] \to [0, \bar{a}] \ \text{ s.t. } \ \alpha_{t,d}(a,b) \leq a.$$

This means that $\alpha_{t,d}(a,b)$ is the investment A makes into the battle at state (t,d) if A's remaining resources are a and B's remaining resources are b. A *pure strategy* for player A is a collection of investment functions for every possible state in $W(1,0)$:

$$\alpha \;=\; \Big\{\alpha_{t,d}(\cdot) : (t,d) \in W(1,0)\Big\}.$$

For player B the definitions are analogous: An investment function at state (t,d) is a mapping

$$\beta_{t,d} : [0,\bar{a}] \times [0,\bar{b}] \to [0,\bar{b}] \;\; \text{s.t.} \;\; \beta_{t,d}(a,b) \leq b,$$

and a pure strategy is a collection

$$\beta \;=\; \Big\{\beta_{t,d}(\cdot) : (t,d) \in W(1,0)\Big\}.$$

Given $(t,d;a,b)$ and profile of strategies (α,β), denote by $\pi^i_{t,d}(\alpha,\beta \,|\, a,b)$ the expected continuation payoff to player $i = A,B$, conditional on having reached information set $(t,d;a,b)$. Thus, the overall expected payoff to player i in the Tug of War is[11]

$$\pi^i(\alpha,\beta) \;=\; \pi^i_{1,0}(\alpha,\beta \,|\, \bar{a},\bar{b}).$$

Strategy α is a *best response to* β *at* $(t,d;a,b)$ if, for every strategy $\alpha' \neq \alpha$, it is true that

$$\pi^A_{t,d}(\alpha,\beta \,|\, a,b) \;\geq\; \pi^A_{t,d}(\alpha',\beta \,|\, a,b).$$

Strategy α is a best response to β at (t,d) if α is a best response to β at $(t,d;a,b)$ $\forall (a,b) \in [0,\bar{a}] \times [0,\bar{b}]$. Player B's best responses are defined similarly. A pair of strategies (α,β) that are mutual best responses at all $(t,d) \in W(1,0)$ is a *pure strategy subgame perfect equilibrium.*

3.3 Remarks

A few remarks are in order. First, our definition of a strategy is Markovian in the following sense: Each "relevant" information set is characterized by a tuple $(t,d;a,b) \in W(1,0) \times [0,\bar{a}] \times [0,\bar{b}]$, consisting of a battle state (t,d) and a pair of remaining budgets (a,b). Conditional on $(t,d;a,b)$, efforts do not depend on the path taken through the

[11] Note that $\pi^i(\cdot)$ is not the probability that i wins the Tug of War, but rather the difference between the probability that i wins and the probability that i loses the Tug of War. These two probabilities are not, generally, complements of each other; see Section 3.3.

state space to reach (t, d). However, conditional on (t, d) only, the actual realized efforts will generally depend on how (t, d) was reached.[12]

Second, since player A wins if and only if a terminal state (t, L) is reached before $(t', -L)$ is reached (and vice versa for B), the game may not result in a winner. For example, consider a strategy profile in which the players take turns investing half of their remaining resources:

$$\alpha_{t,d}(a,b) = \begin{cases} a/2 & \text{if } t \text{ is odd,} \\ 0 & \text{otherwise,} \end{cases} \qquad \beta_{t,d}(a,b) = \begin{cases} b/2 & \text{if } t \text{ is even,} \\ 0 & \text{otherwise.} \end{cases}$$

With these strategies, player A wins the stage battle in every odd period and player B wins the stage battle in every even period, so no terminal state is ever reached (since $L \geq 2$). This example demonstrates that the probability that A wins the game and the probability that B wins the game do not necessarily have to sum to one. We can interpret a situation where neither player wins as a "draw," and our assumptions imply a payoff of zero to both players in the event of a draw. Because of the Blotto budget assumption, our Tug of War game with fixed budgets is a zero-sum game in both cases. All of our results would continue to hold if we assumed that losing the Tug of War and a draw are indifferent outcomes.

4 The One-Step Deviation Principle

The use of recursive methods to solve for the subgame perfect equilibrium of the Tug of War game requires the *one-step deviation principle* to hold: In order to check that a strategy profile is subgame perfect equilibrium, it is sufficient to verify that no player can improve his payoff by switching to a different strategy that deviates from the equilibrium at a single information set and then returns to the equilibrium strategy.

4.1 No continuity at infinity

The "usual way" of establishing that the one-step deviation principle holds is to invoke *continuity at infinity*: For every $\varepsilon > 0$ there exists T such that any two strategy profiles that are identical in the first T periods will differ in their associated payoffs by at most ε. This property is a well-known sufficient condition for the one-step deviation principle (see, e.g., Fudenberg and Tirole 1991, p. 110). It is satisfied, for example, in all repeated games in which players maximize the present value of bounded stage payoffs. Our Tug

[12] The Markovian property only requires that the investment *functions* not be path-dependent. Since the budgets are generally path-dependent even under Markovian strategies, actual realized investments will be path-dependent as well.

of War game, on the other hand, is not continuous at infinity, as the following example demonstrates:

Example 1. Consider the strategy profile

$$\alpha_{t,d}(a,b) = \begin{cases} a/3 & \text{if } t \text{ is odd,} \\ 0 & \text{otherwise,} \end{cases} \qquad \beta_{t,d}(a,b) = \begin{cases} b/2 & \text{if } t \text{ is even,} \\ 0 & \text{otherwise.} \end{cases}$$

With these strategies, player A wins the stage battle in every odd period and player B wins the stage battle in every even period, so that each player's probability of winning the Tug of War is zero, yielding a zero payoff to both.

Now consider the following alternative strategy for player A (while not changing B's strategy): Fix some even $T \geq 2$, use the investment functions $\alpha_{t,d}$ until period T, and then allocate the remaining budget evenly across the next L battles. Formally, this strategy is

$$\alpha'_{t,d}(a,b) = \begin{cases} a/3 & \text{if } t \leq T \text{ is odd,} \\ 0 & \text{if } t \leq T \text{ is even or } t > T+L, \\ a/(T+L+1-t) & \text{if } t \in \{T+1,\dots,T+L\}. \end{cases}$$

Under this strategy, in period $T+1$ the state of the game will be $(T+1, 0)$ with probability 1. Furthermore, if T is large player A's remaining resources at state $(T+1, 0)$ will be approximately $\bar{a}/2$ while B's resources will be approximately zero. If A allocates his remaining resources evenly across the next L battles, he wins each of the next L battles with a high probability. As $T \to \infty$, the probability that A wins the Tug of War with this strategy converges to one, and so does A's expected payoff, while continuity at infinity would have both converge to zero.

We could introduce either a discount factor or some small exogenous probability with which the game breaks down in each period. These assumptions would be mild (and possibly more plausible than our undiscounted payoff structure), and with either modification the game would become continuous at infinity, so that the standard result applies. The problem with this approach is that we would not be able to obtain the closed-form characterization of the equilibrium strategies that we get with our undiscounted specification. Specifically, with discounting the solution to the second-order difference equation (23) in the proof of Lemma 2 would no longer represent a player's continuation payoff.[13]

[13] The solution would still represent the probability that a player wins if he is d victories ahead of his rival and has ρ times as many resources as the rival. However, this is no longer equal to a player's expected present value of winning. To obtain this expected present value, the discount factor would need to be incorporated in differential equation (23), which would considerably complicate its solution.

To avoid these complications down the road, we make the upfront investment of proving that, even though the undiscounted game is not continuous at infinity, the one-step deviation principle still holds.

Definition 1. Consider the profile of strategies $(\alpha, \beta) = \{\alpha_{t,d}(\cdot), \beta_{t,d}(\cdot) : (t,d) \in W(1,0)\}$. Strategy $\widehat{\alpha} = \{\widehat{\alpha}_{t,d}(\cdot) : (t,d) \in W(1,0)\}$ for player A is a *one-step deviation* from α if there exists $(t^*, d^*; a^*, b^*) \in W(1,0) \times [0,\bar{a}] \times [0,\bar{b}]$ such that

$$\alpha_{t,d}(a,b) = \widehat{\alpha}_{t,d}(a,b) \quad \Leftrightarrow \quad (t,d;a,b) \neq (t^*,d^*;a^*,b^*).$$

The one-step deviation is *profitable at the point of deviation* if

$$\pi^A_{t^*,d^*}(\widehat{\alpha}, \beta \,|\, a^*, b^*) > \pi^A_{t^*,d^*}(\alpha, \beta \,|\, a^*, b^*).$$

A profitable one-step deviation for B is defined similarly.

Note that the profitability of a one-step deviation is defined conditional on having reached the information set at which the deviation occurs (reflecting the sequential rationality aspect of subgame perfect equilibrium). We will show:

Theorem 1. *(One-Step Deviation Principle) A pure strategy profile is a subgame perfect equilibrium of the Tug of War with fixed resources if and only if no player has a one-step deviation that is profitable at the point of deviation.*

4.2 Proof of Theorem 1

The "only if" part of Theorem 1 is immediate from the definition of subgame perfect equilibrium. To prove the "if" part, we show that if (α, β) is not a pair of mutual best responses at $(1,0;\bar{a},\bar{b})$ then at least one player has a profitable one-shot deviation. Therefore, we prove the "if" part for Nash equilibrium. Because the same argument can be repeated for any arbitrary subgame $(t,d;a,b)$, by adjusting the notation, the one-step deviation principle for subgame perfect equilibrium is implied.

We first have to introduce a few preliminary definitions (Step 1 below). Thereafter, our argument proceeds in a fashion similar to the proof of the standard result that relies on continuity at infinity: We show that any profitable deviation from the equilibrium must be profitable within a finite number of periods (Step 2). We then show, through backward induction, that this implies the existence of a profitable one-step deviation (Step 3).

Step 1. Without loss of generality we consider only deviations for player A from a candidate equilibrium profile, and our definitions below reflect this choice.

We impose the following complete and strict ordering $\rhd$ on the states in $W(1,0)$:[14]

$$(t', d') \rhd (t, d) \quad \Leftrightarrow \quad \big[t' - d' > t - d\big] \text{ or } \big[t' - d' = t - d \text{ and } d' > d\big].$$

Let $\mathcal{P}(t,d)$ be the immediate successor of (t,d) under the ordering $\rhd$. Likewise, let $\mathcal{P}^{-1}(t,d)$ denote the immediate predecessor of (t,d).

As before, a strategy will be regarded as a collection of investment functions, one for each state $(t,d) \in W(1,0)$. Given two strategies for A, $\alpha = \{\alpha_{t,d}\}$ and $\widetilde{\alpha} = \{\widetilde{\alpha}_{t,d}\}$, and some state (t_0, d_0), define the *blended strategy* $\big\langle \widetilde{\alpha}, \alpha, (t_0, d_0) \big\rangle$ as follows:

$$\big\langle \widetilde{\alpha}, \alpha, (t_0, d_0) \big\rangle = \big\{\alpha_{t,d}^{\text{blended}}\big\}, \quad \text{where} \quad \alpha_{t,d}^{\text{blended}} = \begin{cases} \alpha_{t,d} & \text{if } (t,d) \rhd (t_0, d_0), \\ \widetilde{\alpha}_{t,d} & \text{otherwise.} \end{cases}$$

Note that $(t,d) \in S(t_0, d_0)$ implies $(t,d) \rhd (t_0, d_0)$; hence, blended strategy $\big\langle \widetilde{\alpha}, \alpha, (t_0, d_0) \big\rangle$ is identical to the strategy α at all battle states $(t,d) \in S(t_0, d_0)$. Figure 2 illustrates the ordering $\rhd$ and the construction of blended strategies.

Figure 2: Ordering $\rhd$ and blended strategy $\big\langle \widetilde{\alpha}, \alpha, (4, -1) \big\rangle$ (for $L = 3$).

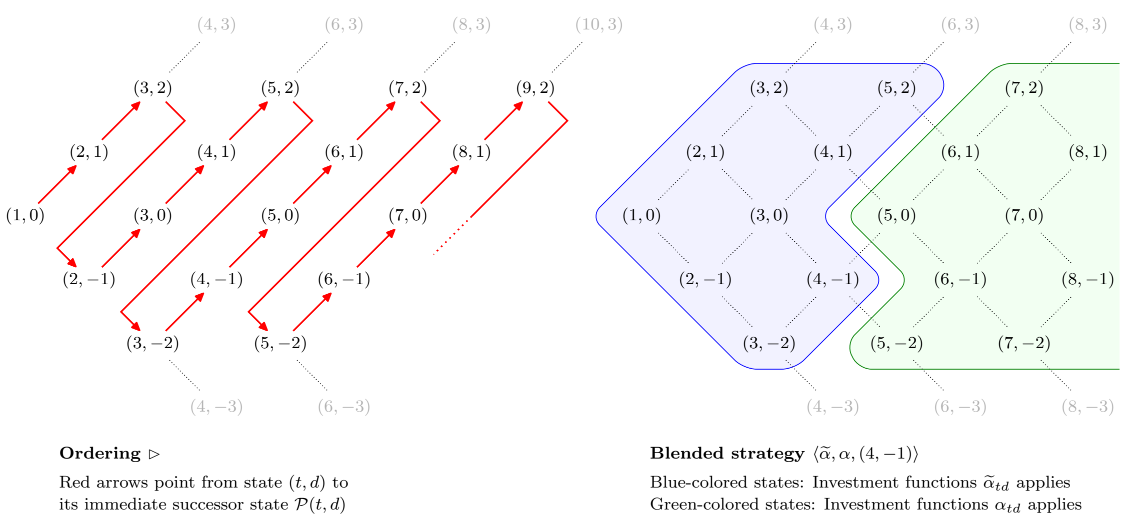

Let $\mathcal{T} = \{L+1, L+3, L+5, \ldots\}$; thus, (t, L) and $(t, -L)$ are terminal states if and only if $t \in \mathcal{T}$. Given $t \in \mathcal{T}$, let $p_t^+(\alpha, \beta)$ be the probability that terminal state (t, L) is

[14] If we considered deviations for player B, we would redefine the ordering to the following:

$$(t', d') \rhd (t, d) \quad \Leftrightarrow \quad \big[t + d' > t + d\big] \text{ or } \big[t' + d' = t + d \text{ and } d' < d\big],$$

so that the red arrows in Figure 2 would be oriented in a south-easterly direction (instead of a north-easterly direction).

reached, and let $p_t^-(\alpha,\beta)$ be the probability that terminal state $(t,-L)$ is reached, under strategy profile (α,β). If we define $p_t(\alpha,\beta) \equiv p_t^+(\alpha,\beta) - p_t^-(\alpha,\beta)$, we can express player A's expected payoff as

$$\pi^A(\alpha,\beta) \;=\; \sum_{t\in\mathcal{T}} p_t(\alpha,\beta). \tag{1}$$

The following observation can be verified by inspecting the ordering $\rhd$, as illustrated graphically on the left side of Figure 2: Fix two battle states (t_0,d_0) and (t_1,d_1), with $(t_1,d_1)\rhd(t_0,d_0)$. Note that for any terminal state (t,d) it must be true that $t > t_0 - d_0 - L$. Since blended strategy $\langle\widetilde{\alpha},\alpha,(t_0,d_0)\rangle$ differs from strategy $\widetilde{\alpha}$ only at states $(t_1,d_1)\rhd(t_0,d_0)$, this implies that

$$\forall t\in\mathcal{T} \text{ s.t. } t \le t_0 - d_0 - L: \quad p_t\big(\langle\widetilde{\alpha},\alpha,(t_0,d_0)\rangle,\,\sigma^B\big) = p_t(\widetilde{\alpha},\alpha). \tag{2}$$

Step 2. Suppose the profile (α,β) is not a subgame perfect equilibrium. Focusing on player A without loss of generality, there exists a strategy $\widetilde{\alpha}$, a state $(t,d)\in W(1,0)$, and a set of budgets $(a,b)\in[0,\bar{a}]\times[0,\bar{b}]$ such that

$$\pi^A_{t,d}(\widetilde{\alpha},\beta\,|\,a,b) \;>\; \pi^A_{t,d}(\alpha,\beta\,|\,a,b). \tag{3}$$

As discussed earlier, we may also assume that $(t,d)=(1,0)$ and $(a,b)=(\bar{a},\bar{b})$. Using (1), (3) then becomes

$$\pi^A(\widetilde{\alpha},\beta) \;=\; \sum_{t\in\mathcal{T}} p_t(\widetilde{\alpha},\beta) \;>\; \pi^A(\alpha,\beta).$$

It follows that there exists $T\in\mathcal{T}$ such that

$$\sum_{t\in\mathcal{T},\,t\le T} p_t(\widetilde{\alpha},\beta) \;>\; \pi^A(\alpha,\beta). \tag{4}$$

Fix any such T and consider the blended strategy $\langle\widetilde{\alpha},\alpha,(T+L,0)\rangle$. Player A's expected payoff when using this blended strategy against strategy σ^B is

$$\begin{aligned}
\pi^A\big(\langle\widetilde{\alpha},\alpha,(T+L,0)\rangle,\,\beta\big) \;&=\; \sum_{t\in\mathcal{T}} p_t\big(\langle\widetilde{\alpha},\alpha,(T+L,0)\rangle,\,\beta\big)\\
&=\; \sum_{t\in\mathcal{T},\,t\le T} p_t(\widetilde{\alpha},\beta) \;+\; \sum_{t\in\mathcal{T},\,t>T} p_t\big(\langle\widetilde{\alpha},\alpha,(T+L,0)\rangle,\,\beta\big)\\
&\ge\; \sum_{t\in\mathcal{T},\,t\le T} p_t(\widetilde{\alpha},\beta) \;>\; \pi^A(\alpha,\beta),
\end{aligned}$$

where the second equality is by (2) and the final inequality is by (4). Thus, if in profile (α, β) player A has a profitable deviation, then player A has a profitable deviation that differs from strategy α at a finite number of states (namely, a blended strategy consisting of the deviation and the original strategy α).

Step 3. Consider player A's investment decision at state $(T+L, 0)$. Suppose there exists $(a^*, b^*) \in [0, \bar{a}] \times [0, \bar{b}]$ such that

$$\pi^A_{T+L,0}\big(\langle \widetilde{\alpha}, \alpha, (T+L, 0)\rangle, \beta \,\big|\, a^*, b^*\big) \;>\; \pi^A_{T+L,0}\big(\alpha, \beta \,\big|\, a^*, b^*\big).$$

Since the blended strategy $\langle \widetilde{\alpha}, \alpha, (T+L, 0)\rangle$ prescribes the same investment functions as strategy as α at all $(t, d) \in S(T+L, 0)$ (see Step 1), the strategy $\widehat{\alpha} = \big\{\widehat{\alpha}_{t,d}(\cdot)\big\}$ with

$$\widehat{\alpha}_{t,d}(a, b) = \begin{cases} \widetilde{\alpha}_{t,d}(a, b) & \text{if } (t, d; a, b) = (T+L, 0; a^*, b^*), \\ \alpha_{t,d}(a, b) & \text{otherwise} \end{cases}$$

is a profitable one-step deviation for player A. In this case, the proof is complete.

If, on the other hand, no such (a^*, b^*) exists, we proceed as follows. Since player A does not gain at any $(T+L, 0; a, b)$ from playing blended strategy $\langle \widetilde{\alpha}, \alpha, (T+L, 0)\rangle$, despite the fact that $\pi^A\big(\langle \widetilde{\alpha}, \alpha, (T+L, 0)\rangle, \beta\big) \;>\; \pi^A(\alpha, \beta)$ (as shown in Step 2), it must be true that

$$\pi^A\big(\langle \widetilde{\alpha}, \alpha, \mathcal{P}^{-1}(T+L, 0)\rangle, \beta\big) \;>\; \pi^A(\alpha, \beta).$$

If there exists $(a^*, b^*) \in [0, \bar{a}] \times [0, \bar{b}]$ such that

$$\pi^A_{\mathcal{P}^{-1}(T+L,0)}\big(\langle \widetilde{\alpha}, \alpha, \mathcal{P}^{-1}(T+L, 0)\rangle, \beta \,\big|\, a^*, b^*\big) \;>\; \pi^A_{\mathcal{P}^{-1}(T+L,0)}\big(\alpha, \beta \,\big|\, a^*, b^*\big),$$

one can construct a profitable one-step deviation by using the same construction as above. In this case, the proof is complete. If, on the other hand, no such (a^*, b^*) exists, it must be true that

$$\pi^A\big(\langle \widetilde{\alpha}, \alpha, \mathcal{P}^{-1}(\mathcal{P}^{-1}(T+L, 0))\rangle, \beta\big) \;>\; \pi^A(\alpha, \beta),$$

and so on.

This process is guaranteed to yield a profitable one-step deviation. Note that, eventually, $\mathcal{P}^{-1}(\mathcal{P}^{-1}(\mathcal{P}^{-1} \;\ldots\; \mathcal{P}^{-1}(T+L, 0))) = (1, 0)$, and if no profitable one-step deviation has been found until then, it must be true that

$$\pi^A\big(\langle \widetilde{\alpha}, \alpha, (1, 0)\rangle, \beta\big) \;>\; \pi^A(\alpha, \beta).$$

But this is immediately implies that the strategy $\widehat{\alpha} = \left\{\widehat{\alpha}_{t,d}(\cdot)\right\}$ with

$$\widehat{\alpha}_{t,d}(a,b) = \begin{cases} \widetilde{\alpha}_{t,d}(\bar{a},\bar{b}) & \text{if } (t,d;a,b) = (1,0;\bar{a},\bar{b}), \\ \alpha_{t,d}(a,b) & \text{otherwise} \end{cases}$$

is a profitable one-step deviation. □

5 Derivation of the Equilibrium Strategy Profile

For the remainder of the paper, we assume that the contest success function is

$$f(x,y) \;=\; \begin{cases} \dfrac{x^\gamma}{x^\gamma + y^\gamma} & \text{if } x+y>0, \\ 1/2 & \text{if } x+y=0, \end{cases} \tag{5}$$

with $\gamma > 0$. The expression in (5) is the *generalized Tullock function* (see, e.g., Tullock 1980; Baye, Kovenock, and De Vries 1994), with the parameter γ measuring the responsiveness of a player's success probability to changes in the player's own effort.

Using this class of contest success functions, we analytically derive a candidate profile for pure strategy, subgame perfect equilibrium in the Tug of War game with fixed resources. The profile is a candidate insofar as its derivation is based on the players' first-order conditions for payoff maxima only. We will go on to verify numerically that a higher-order conditions are also satisfied, provided the parameter γ is not too large.

We can define four properties describing various aspects of "self-similarity" of strategy profiles (α, β). Our candidate equilibrium will satisfy all.

Definition 2. (i) (α, β) is *stationary* if for all (t,d) and (t',d):

$$\alpha_{t,d}(a,b) = \alpha_{t',d}(a,b) \;\text{ and }\; \beta_{t,d}(a,b) = \beta_{t',d}(a,b) \quad \forall (a,b).$$

(ii) (α, β) is *symmetric* if for all (t,d):

$$\alpha_{t,d}(a,b) = \beta_{t,-d}(b,a) \quad \forall (a,b).$$

(iii) (α, β) is *homogeneous* if for all (t,d):

$$\alpha_{t,d}(\lambda a, \lambda b) = \lambda \alpha_{t,d}(a,b) \;\text{ and }\; \beta_{t,d}(\lambda a, \lambda b) = \lambda \beta_{t,d}(a,b) \quad \forall (a,b),\ \forall \lambda > 0.$$

(iv) (α, β) is *balanced* if for all (t,d):

$$\alpha_{t,d}(a,b)/a = \beta_{t,d}(a,b)/b \quad \forall (a,b) \gg 0.$$

Note that all four properties are independent, and none implies any of the others. Property (i) states that two subgames that differ only in t exhibit identical continuation behavior. This property reduces the number of battle states that have to be considered in the analysis from infinity to $2L-1$. Properties (ii)–(iv) are directly related to the mathematical properties of our contest success function. Property (ii) states that in a subgame that is a "mirror image" of another, behavior should be a "mirror image" of behavior in the other. This is an expected property if the contest success function is symmetric. Property (iii) states that the shares of their remaining resources that players invests in a given battle depend on their relative, but not absolute, budgets. This is a sensible property if the contest success function is homogenous of degree zero, as it is here. Finally, property (iv) states that, at each battle, the two players invest identical fractions of their remaining budgets; thus, their relative budgets remain unchanged throughout the contest. We will show that this property, too, is generated by the homogeneity of f and implies that a player's probability of winning each battle remains constant.

We now construct a candidate equilibrium in three steps. In the first step, we show the following: If a pure strategy profile is a subgame perfect equilibrium and both players use balanced strategies starting at some time $t+1$, then they must use balanced strategies also at time t. This step allows us to search for an equilibrium in which the players invest identical fractions of their remaining resources in every battle. In the second step, we will find the common fraction of their budgets that the players invest, based on players' first-order conditions for profit maxima. This fraction will still depend on d, a, and b, but not on the player. In the third step, we verify higher-order conditions computationally.

5.1 Step 1: Optimality of balanced strategies

Take a strategy profile (α, β). Let $\pi_{t,d}(a,b)$ be player A's expected payoff conditional on the players entering state (t,d) with resources (a,b) remaining. If (t,d) is a terminal state, then $\pi_{t,d}(a,b) = 1$ if $d = L$ and $\pi_{t,d}(a,b) = -1$ if $d = -L$.

Now fix a battle state (t,d). Suppose that, for all states (t',d') with $t' = t+1$, the profile (α,β) is a balanced continuation Nash equilibrium. We will show the following: If (α,β) is a continuation Nash equilibrium also at (t,d), then it is balanced.

The two immediate successor states of (t,d) are $(t+1,d+1)$ and $(t+1,d-1)$. Since t and d are fixed, we save on notation by defining

$$V(a,b) \;=\; \pi_{t+1,d+1}(a,b) \quad \text{and} \quad W(a,b) \;=\; \pi_{t+1,d-1}(a,b).$$

Fix a pair of budgets (a,b) at state (t,d). If player A invests x in state (t,d), he will have resources $a-x$ in each of the two successor states. Similarly, if B invests y in state (t,d), he will have $b-y$ in each of the two successor states. Since (α,β) is a continuation equilibria at $(t+1, d+1)$ and $(t+1, d-1)$, player A solves

$$\max_{0 \le x \le a} \quad f(x,y)V(a-x, b-y) + \bigl(1 - f(x,y)\bigr)W(a-x, b-y).$$

The first-order condition is

$$\begin{aligned} \frac{\gamma x^{\gamma-1}y^{\gamma}}{(x^{\gamma}+y^{\gamma})^2}\Bigl[V(a-x,b-y) - W(a-x,b-y)\Bigr] \\ = \frac{x^{\gamma}}{x^{\gamma}+y^{\gamma}}V_a(a-x,b-y) + \frac{y^{\gamma}}{x^{\gamma}+y^{\gamma}}W_a(a-x,b-y). \end{aligned} \tag{6}$$

Since the profile (α, β) is balanced at states $(t+1, d+1)$ and $(t+1, d-1)$ by assumption, and since f is homogeneous of degree zero, V and W must be homogeneous of degree zero. Thus, if we define

$$v(a/b) \;\equiv\; V(a/b, 1) = \pi_{t+1,d+1}(a/b, 1) \quad \text{and} \quad w(a/b) \;\equiv\; W(a/b, 1) = \pi_{t+1,d-1}(a/b, 1)$$

then (6) can be expressed as

$$\begin{aligned} \frac{\gamma x^{\gamma-1}y^{\gamma}}{(x^{\gamma}+y^{\gamma})^2}\left[v\left(\frac{a-x}{b-y}\right) - w\left(\frac{a-x}{b-y}\right)\right] \\ = \frac{1}{b-y}\left[\frac{x^{\gamma}}{x^{\gamma}+y^{\gamma}}v'\left(\frac{a-x}{b-y}\right) + \frac{y^{\gamma}}{x^{\gamma}+y^{\gamma}}w'\left(\frac{a-x}{b-y}\right)\right]. \end{aligned} \tag{7}$$

Now turn to player B. This player solves

$$\min_{0 \le y \le b} \quad f(x,y)V(a-x, b-y) + \bigl(1 - f(x,y)\bigr)W(a-x, b-y).$$

The first-order condition is

$$\begin{aligned} -\frac{\gamma x^{\gamma}y^{\gamma}}{(x^{\gamma}+y^{\gamma})^2}\Bigl[V(a-x,b-y) - W(a-x,b-y)\Bigr] \\ = \frac{x^{\gamma}}{x^{\gamma}+y^{\gamma}}V_b(a-x,b-y) + \frac{y^{\gamma}}{x^{\gamma}+y^{\gamma}}W_b(a-x,b-y), \end{aligned} \tag{8}$$

and making the same substitutions as before, we can express (8) as

$$\begin{aligned} -\frac{\gamma x^{\gamma-1}y^{\gamma}}{(x^{\gamma}+y^{\gamma})^2}\left[v\left(\frac{a-x}{b-y}\right) - w\left(\frac{a-x}{b-y}\right)\right] \\ = -\frac{a-x}{(b-y)^2}\left[\frac{x^{\gamma}}{x^{\gamma}+y^{\gamma}}v'\left(\frac{a-x}{b-y}\right) + \frac{y^{\gamma}}{x^{\gamma}+y^{\gamma}}w'\left(\frac{a-x}{b-y}\right)\right]. \end{aligned} \tag{9}$$

Dividing (9) by (7) we get the following condition:

$$\frac{x}{y} = \frac{a-x}{b-y} \quad \Rightarrow \quad \frac{x}{a} = \frac{y}{b}. \tag{10}$$

That is, the two players invest identical fractions of their remaining budgets in the battle at state (t, d). The same argument can be made for all $(a, b) \in [0, \bar{a}] \times [0, \bar{b}]$, and for all (t', d') with $t' = t$. It follows that the profile (α, β) is balanced at (t, d') for all d'.

If players invest identical fractions of their budgets, then in the next state their relative remaining budgets are exactly the same as in the previous state. Thus, starting at any state with relative budgets a/b, player A wins every subsequent battle with probability $a^\gamma/(a^\gamma + b^\gamma)$. We can show that this implies:

Lemma 2. *Suppose (α, β) is a profile of balanced strategies. Then at any state (t, d), the equilibrium probability that A wins the Tug of War starting at the subgame game $(t, d; a, b)$ depends only on d and $\rho = a/b$, and is given by*

$$u_d(\rho) = \begin{cases} \dfrac{\rho^{\gamma(L-d)} - \rho^{2\gamma L}}{1 - \rho^{2\gamma L}} & \text{if } \rho \neq 1, \\ (L+d)/(2L) & \text{if } \rho = 1. \end{cases} \tag{11}$$

By switching roles, player B's probability of winning in the same subgame must be $u_{-d}(\rho^{-1})$. It is easy to verify that $u_d(\rho) + u_{-d}(\rho^{-1}) = 1$; thus, when both players use balanced strategies the probability of a draw is zero.[15] Because there is no discounting, we can write A's continuation utility if he wins the battle at state (t, d) and is left with relative budget ρ in the next battle as

$$u_{d+1}(\rho) + (1 - u_{d+1}(\rho))(-1) = 2u_{d+1}(\rho) - 1.$$

Simililary, if he loses the battle at state (t, d) and is left with relative budget ρ in the next battle, his continuation utility is

$$u_{d-1}(\rho) + (1 - u_{d-1}(\rho))(-1) = 2u_{d-1}(\rho) - 1.$$

[15]Fundamentally, this is a consequence of Kolmogorov's zero-one law (see, e.g., Feller 1970, p. 124). Given a balanced profile of strategies, a draw is a tail event of the *statistically independent* outcomes of the individual battles. The probability of such tail events is either zero or one. It cannot be one, however, because each player clearly has some positive probability of winning (e.g., the probability that A wins each of the first L battles in a balanced profile is $\left[a^\gamma/(a^\gamma + b^\gamma)\right]^L > 0$). Thus, a draw must have a zero probability.

Therefore, we can without loss of generality assume that players choose their investments to maximize their probability of winning. In the following, we accordingly redefine $v(\rho) := u_{d+1}(\rho)$ and $w(\rho) := u_{d-1}(\rho)$.

5.2 Step 2: Pinning down the investment functions

Any balanced strategy profile can be described as a set of functions

$$\Big\{ s_{t,d} : [0,\bar{a}] \times [0,\bar{b}] \to [0,1] : (t,d) \in S(1,0) \Big\},$$

which has the interpretation that, if state (t,d) is reached and the remaining budgets are (a,b), player A invests $\alpha_{t,d}(a,b) = s_{t,d}(a,b)a$ and player B invests $\beta_{t,d}(a,b) = s_{t,d}(a,b)b$. To find this function in equilibrium at state (t,d), we can use the conditions $x = sa$ and $y = sb$ in the first-order condition of either player A or player B.

Substituting $x = sa$ and $y = sb$ into A's first-order condition (7), we obtain

$$\frac{\gamma a^{\gamma-1} b^{\gamma}}{s(a^{\gamma}+b^{\gamma})^2}\left[v\left(\frac{a}{b}\right) - w\left(\frac{a}{b}\right)\right] \;=\; \frac{1}{(1-s)b}\left[\frac{a^{\gamma}}{a^{\gamma}+b^{\gamma}} v'\left(\frac{a}{b}\right) + \frac{b^{\gamma}}{a^{\gamma}+b^{\gamma}} w'\left(\frac{a}{b}\right)\right]. \tag{12}$$

Now define $\rho \equiv a/b$ and rearrange (12) to

$$\frac{1-s}{s}\frac{\gamma\rho^{\gamma-1}}{1+\rho^{\gamma}}\big[v(\rho) - w(\rho)\big] \;=\; \rho^{\gamma} v'(\rho) + w'(\rho). \tag{13}$$

where $\rho \equiv a/b$. It is clear from (13) that, if the equilibrium is balanced, it must also be homogeneous: s depends only on the ratio $\rho = a/b$, but not on a and b individually.

Assume for the time being that $\rho \neq 1$. Then, using the first term in (11), we can write

$$v(\rho) \;=\; u_{d+1}(\rho) \;=\; \frac{\rho^{\gamma(L-d-1)} - \rho^{2\gamma L}}{1-\rho^{2\gamma L}}$$

and

$$w(\rho) \;=\; u_{d-1}(\rho) \;=\; \frac{\rho^{\gamma(L-d+1)} - \rho^{2\gamma L}}{1-\rho^{2\gamma L}}.$$

Thus, the left-hand side of (13) becomes

$$\frac{1-s}{s}\frac{\gamma\rho^{\gamma-1}}{1+\rho^{\gamma}}\big[v(\rho) - w(\rho)\big] \;=\; \frac{1-s}{s}\frac{\gamma\rho^{\gamma-1}}{1+\rho^{\gamma}}\frac{\rho^{\gamma(L-d-1)}(1-\rho^{2\gamma})}{1-\rho^{2\gamma L}}. \tag{14}$$

The derivatives of v and w are

$$v'(\rho) \;=\; \frac{\gamma(L-d-1)\rho^{\gamma(L-d-1)-1} - 2\gamma L\rho^{2\gamma L-1}}{1-\rho^{2\gamma L}} \;+\; \frac{\Big[\rho^{\gamma(L-d-1)} - \rho^{2\gamma L}\Big]2\gamma L\rho^{2\gamma L-1}}{(1-\rho^{2\gamma L})^2}$$

$$= \gamma \frac{\rho^{\gamma(L-d-1)-1}}{1-\rho^{2\gamma L}} \Big[L-d-1-2L\rho^{\gamma(L+d+1)} \Big] + 2\gamma L \frac{\rho^{2\gamma L-1}}{1-\rho^{2\gamma L}} v(\rho)$$

and

$$w'(\rho) = \frac{\gamma(L-d+1)\rho^{\gamma(L-d+1)-1} - 2\gamma L \rho^{2\gamma L-1}}{1-\rho^{2\gamma L}} + \frac{\Big[\rho^{\gamma(L-d+1)} - \rho^{2\gamma L}\Big] 2\gamma L \rho^{2\gamma L-1}}{(1-\rho^{2\gamma L})^2}$$

$$= \gamma \frac{\rho^{\gamma(L-d+1)-1}}{1-\rho^{2\gamma L}} \Big[L-d+1-2L\rho^{\gamma(L+d-1)} \Big] + 2\gamma L \frac{\rho^{2\gamma L-1}}{1-\rho^{2\gamma L}} w(\rho).$$

Thus, the right-hand side of (13) becomes

$$\rho^\gamma v'(\rho) + w'(\rho) = \gamma \frac{\rho^{\gamma(L-d)-1}}{1-\rho^{2\gamma L}} \Big[(L-d-1) + \rho^\gamma (L-d+1) - (1+\rho^\gamma) 2L \rho^{\gamma(L+d)} \Big] + 2\gamma L \frac{\rho^{2\gamma L-1}}{1-\rho^{2\gamma L}} \Big[\rho^\gamma v(\rho) + w(\rho) \Big]. \quad (15)$$

Since we can always express

$$\frac{\rho^\gamma}{1+\rho^\gamma} v(\rho) + \frac{1}{1+\rho^\gamma} w(\rho) = \pi_d(\rho) \left(= \frac{\rho^{\gamma(L-d)} - \rho^{2\gamma L}}{1-\rho^{2\gamma L}} \right),$$

we can further rewrite (15) as follows:

$$\begin{aligned}
\rho^\gamma v'(\rho) + w'(\rho) &= \gamma \frac{\rho^{\gamma(L-d)-1}}{1-\rho^{2\gamma L}} \Big[(L-d-1) + \rho^\gamma (L-d+1) - (1+\rho^\gamma) 2L \rho^{\gamma(L+d)} \Big] \\
&\quad + (1+\rho^\gamma) 2\gamma L \frac{\rho^{2\gamma L-1}}{1-\rho^{2\gamma L}} \, \frac{\rho^{\gamma(L-d)} - \rho^{2\gamma L}}{1-\rho^{2\gamma L}} \\
&= \gamma \frac{\rho^{\gamma(L-d)-1}}{1-\rho^{2\gamma L}} \Big[(L-d-1) + \rho^\gamma (L-d+1) - (1+\rho^\gamma) 2L \rho^{\gamma(L+d)} \\
&\quad + (1+\rho^\gamma) 2L \rho^{\gamma(L+d)} \frac{\rho^{\gamma(L-d)} - \rho^{2\gamma L}}{1-\rho^{2\gamma L}} \Big] \\
&= \gamma \frac{\rho^{\gamma(L-d)-1}}{1-\rho^{2\gamma L}} \Big[(1+\rho^\gamma)(L-d) - (1-\rho^\gamma) \\
&\quad - (1+\rho^\gamma) 2L \rho^{\gamma(L+d)} \left(1 - \frac{\rho^{\gamma(L-d)} - \rho^{2\gamma L}}{1-\rho^{2\gamma L}} \right) \Big] \\
&= \gamma \frac{\rho^{\gamma(L-d)-1}}{1-\rho^{2\gamma L}} \Big[(1+\rho^\gamma)(L-d) - (1-\rho^\gamma) \\
&\quad - (1+\rho^\gamma) 2L \frac{\rho^{\gamma(L+d)} - \rho^{2\gamma L}}{1-\rho^{2\gamma L}} \Big]
\end{aligned}$$

$$= \gamma \frac{\rho^{\gamma(L-d)-1}}{1-\rho^{2\gamma L}} \Big[(1+\rho^{\gamma}) L \Big(1 - 2 \frac{\rho^{\gamma(L+d)} - \rho^{2\gamma L}}{1-\rho^{2\gamma L}} \Big) - (1+\rho^{\gamma}) d - (1-\rho^{\gamma}) \Big]. \quad (16)$$

Plugging (14) and (16) into (13), A's first-order condition becomes

$$\frac{1-s}{s} \frac{\gamma \rho^{\gamma-1}}{1+\rho^{\gamma}} \frac{\rho^{\gamma(L-d-1)}(1-\rho^{2\gamma})}{1-\rho^{2\gamma L}} = \gamma \frac{\rho^{\gamma(L-d)-1}}{1-\rho^{2\gamma L}} \times \Big[(1+\rho^{\gamma}) L \Big(1 - 2 \frac{\rho^{\gamma(L+d)} - \rho^{2\gamma L}}{1-\rho^{2\gamma L}} \Big) - (1+\rho^{\gamma}) d - (1-\rho^{\gamma}) \Big].$$

After cancelling common terms and rearranging, this becomes

$$\frac{1-s}{s} = \frac{1+\rho^{\gamma}}{1-\rho^{\gamma}} L \Big(1 - 2 \frac{\rho^{\gamma(L+d)} - \rho^{2\gamma L}}{1-\rho^{2\gamma L}} - d \Big) - 1, \quad (17)$$

and solving (17) for s and simplifying, we obtain the following expression for the resource share invested by both players:

$$s^* = \frac{1-\rho^{\gamma}}{1+\rho^{\gamma}} \left[2L \frac{1-\rho^{\gamma(L+d)}}{1-\rho^{2\gamma L}} - L - d \right]^{-1}. \quad (18)$$

That is, if player B invests $y = s^* b$ in the current battle, it is optimal for A to invest $x = s^* a$ in the current battle, assuming that investments in all future battles are balanced. (The same solution can be obtained from B's first-order condition.)

If $\rho = 1$ then (18) cannot be evaluated. Using the second term in (11), we have

$$v(1) = \frac{L+d+1}{2L} \quad \text{and} \quad w(1) = \frac{L+d-1}{2L}.$$

In the Appendix, we verify that (11) is continuous and differentiable at $\rho = 1$, with

$$u'_d(1) = \gamma \frac{(L+d)(L-d)}{4L}.$$

Thus,

$$v'(1) = \gamma \frac{(L+d+1)(L-d-1)}{4L} \quad \text{and} \quad w'(1) = \gamma \frac{(L+d-1)(L-d+1)}{4L}.$$

Plugging these terms for v, w, v', and w' into (13) and solving for s, we obtain

$$s^* = \frac{1}{(L+d)(L-d)}. \quad (19)$$

Finally, combining (18) and (19), the candidate equilibrium is described by a family of $2L-1$ functions $s_d^*: (0,\infty) \to [0,1]$, for $d = -L+1, \ldots, L-1$:

$$s_d^*(\rho) \;=\; \begin{cases} \dfrac{1-\rho^\gamma}{1+\rho^\gamma}\left[2L\dfrac{1-\rho^{\gamma(L+d)}}{1-\rho^{2\gamma L}} - L - d\right]^{-1} & \text{if } \rho \neq 1, \\[2ex] \dfrac{1}{(L+d)(L-d)} & \text{if } \rho = 1. \end{cases} \tag{20}$$

The equilibrium investment functions at state (t,d) and budgets $(a,b) \gg (0,0)$ are then constructed as follows:

$$\alpha_{t,d}^*(a,b) \;=\; s_d^*(a/b)\cdot a \quad \text{and} \quad \beta_{t,d}^*(a,b) \;=\; s_d^*(a/b)\cdot b.$$

If one of the players has a zero budget remaining, and the other has a positive budget, then $\rho = 0$ or $\rho = \infty$. The player with the zero budget has no choice but to invest zero in each remaining battle. Thus, the player whose remaining budget is positive is guaranteed to win the contest if he invests any positive amount in the next $L-d$ battles.

5.3 Step 3: Numerical verification

Consider a situation where the Tug of War contest is in state (t,d), the remaining budgets are a and b, and player B invests $s_d^*(a/b)$ of his remaining budget b in the current battle. Assuming that the strategies from period $t+1$ onward are balanced, A' payoff from investing share s of his own budget, a, can be written as

$$\begin{aligned} U_d(s|\rho) \;\equiv\; & \frac{(s\rho)^\gamma}{(s\rho)^\gamma + s_d^*(\rho)^\gamma} u_{d+1}\left(\frac{(1-s)\rho}{1-s_d^*(\rho)}\right) \\ & + \frac{s_d^*(\rho)^\gamma}{(s\rho)^\gamma + (s_d^*(\rho))^\gamma} u_{d-1}\left(\frac{(1-s)\rho}{1-s_d^*(\rho)}\right), \end{aligned} \tag{21}$$

where $\rho = a/b$, $s_d^*(\cdot)$ is given in (20), and $u_d(\cdot)$ is given in (11). An analogous expression for B's payoff at the same decision can be obtained by switching d to $-d$ and ρ to $1/\rho$ in expression (21); hence, it is without loss of generality to focus only on player A. We have to ensure, therefore, that the expression in (21) is always maximized at $s = s_d^*(\rho)$.

To motivate our verification procedure, we begin by discussing several illustrative examples. Figure 3 plots $U_D(s|\rho)$ for $L = 6$, at nine decisions where player A is either -3, 0, or $+3$ victories ahead of player B (i.e., $d \in \{-3, 0, 3\}$), and where player A has either half of B's resources, the same resources as B, or two times B's resources (i.e., $\rho = \{0.5, 1, 2\}$). The parameter γ in the contest success function (5) is set equal to one. In all nine cases, A's payoff is maximized when $s = s_d^*(\rho)$. Thus, in the cases shown in

Figure 3 the first-order conditions do, in fact, imply payoff maxima for A. For any of the depicted combinations of d and ρ, player B's payoff function is identical to what player A's payoff function would be in the "opposing" combination, $-d$ and $1/\rho$. Thus, B's payoff is maximized at the same fractional investment that is optimal for A. Therefore, in the cases shown in Figure 3, the investment share given in (20) is a mutual best reply, provided both players follow this strategy at all subsequent decisions.

The same is no longer true in the cases shown in Figure 4. This figure depicts payoff function $U_d(s|\rho)$ for $L = 6$, at three decision where A is one battle victory away from

Figure 3: Examples of "well behaved" payoff functions ($L = 6$, $\gamma = 1$).

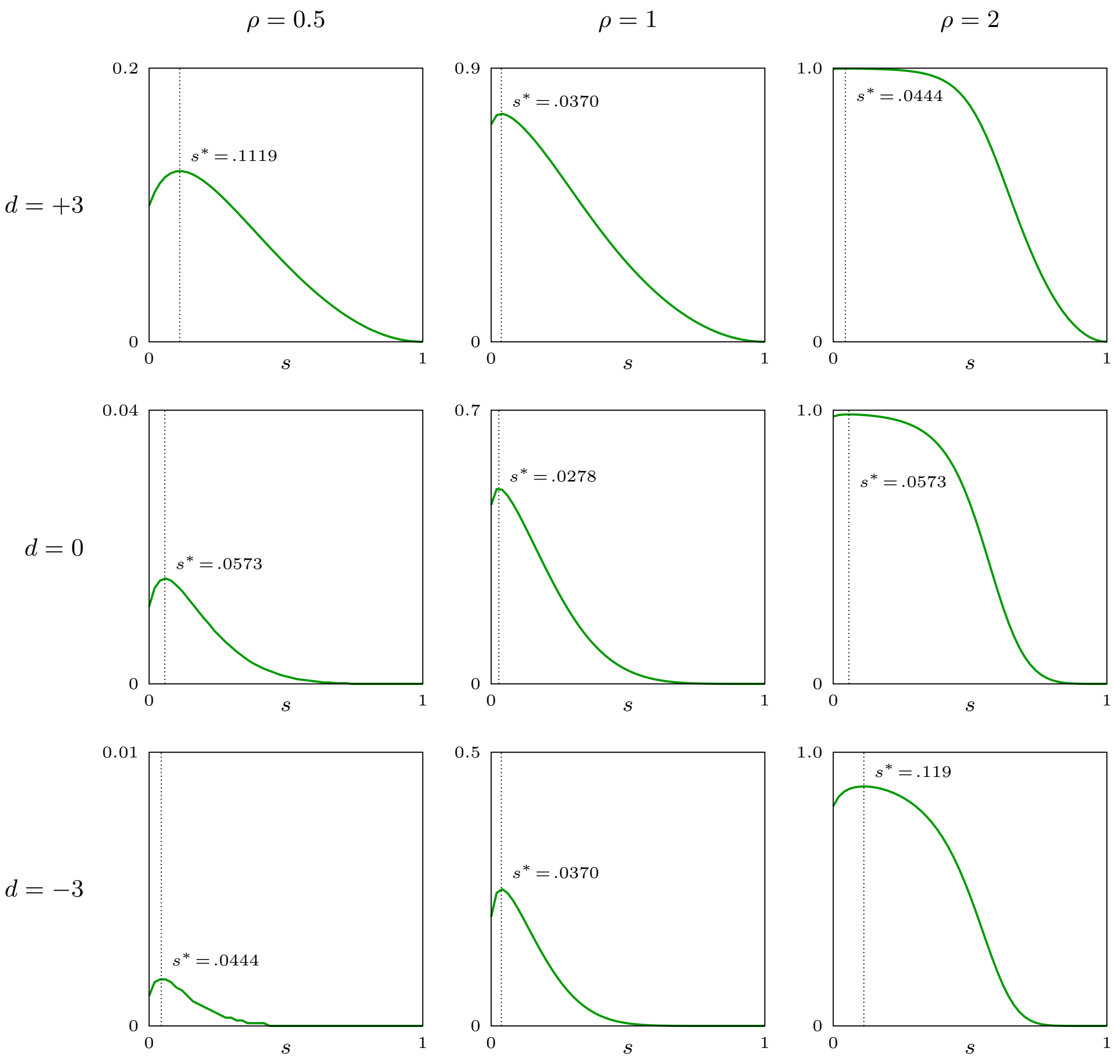

Notes: Graphs show the function $U_d(s|\rho)$ given in (21), for $L = 6$ and $\gamma = 1$. The variable ρ denotes relative resources a/b. Payoffs are shown as a function of $s = x/a$, assuming $y/b = s^*_d(\rho)$, where $s^*_d(\rho)$ is given in (20). In all cases depicted, player A's payoff maximum occurs at $x/a = s^*_d(\rho)$.

winning the game ($d = +5$) and has a smaller remaining effort budget than player B ($\rho = a/b \in \{0.2, 0.5, 0.9\}$). In the panels in the top row, the parameter γ in the contest success function is equal to one. In all three cases, $s = s_d^*(\rho)$ does not maximize A's payoff: In the panel on the left it yields a payoff minimum, with the maximum occurring at $s = 0$; and in the center panel it yields a payoff minimum, with the maximum occurring at $s = 1$. In the right panel, investing $x/a = s_d^*(\rho)$ yields a local payoff maximum, with the global maximum occurring at $s = 1$. Since subgames with $d = 5$ and $\rho \in \{0.2, 0.5, 0.9\}$ could be reached (possibly after deviations), these examples are sufficient to demonstrate that our candidate profile is not a subgame perfect equilibrium when $\gamma = 1$.

Figure 4 also shows how the payoff function changes when the value of γ is slightly lowered, to $\gamma = 0.98, 0.96, 0.94$. In all three cases shown, the payoff gain A could earn when deviating from the prescribed investment $s_d^*(\rho)$ decreases. This suggests that, if γ becomes small enough, these deviations are no longer profitable. Indeed, the panels in the bottom row of Figure 4 plot the function $U_d(s|\rho)$ in the same scenarios, except that $\gamma = 0.5$ (instead of $\gamma = 1$). The payoff maxima now occur exactly at $s_d^*(\rho)$; that is, the problematic cases shown in the top row have now become well behaved. Thus, making

Figure 4: Examples of "badly behaved" payoff functions ($L = 6$, $d = 5$).

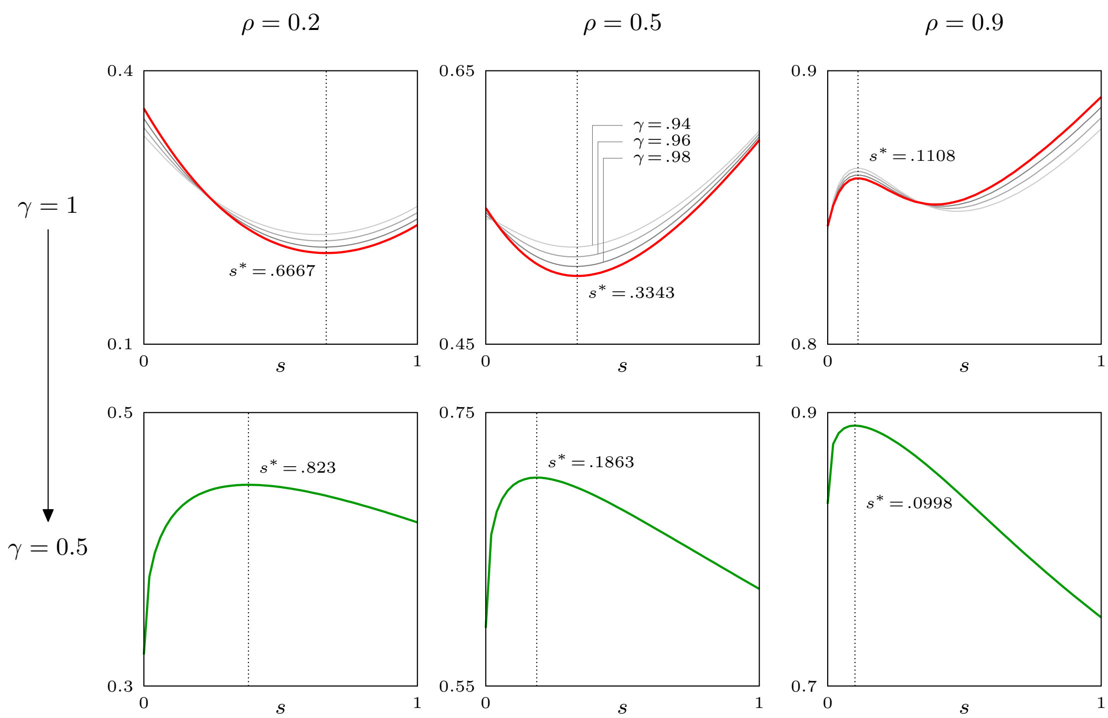

Notes: Graphs show the function $U_d(s|\rho)$ given in (21), for $L = 6$ and $d = 5$. The variable ρ denotes relative resources a/b. Payoffs are shown as a function of $s = x/a$, assuming $y/b = s_d^*(\rho)$, where $s_d^*(\rho)$ is given in (20). When γ is large, $s = s_d^*(\rho)$ does not maximize $U_d(s|\rho)$.

the contest success function that governs each battle "more concave" in a player's own effort also makes a player's payoff function at any decision node "more concave." This, in turn, makes it more likely that an investment that satisfies a player's first order condition is a global payoff maximum.

Motivated by the preceding examples, our objective is to identify a range of values for the parameter γ for which the investments prescribed by (20) are mutual best responses for every $d = -L+1, \ldots, L-1$ and at all $\rho \in (0, \infty)$. As argued above, this is equivalent to showing that $U_d(s|\rho)$ is maximized at $s = s_d^*(\rho)$. There are infinitely many values in the domain of ρ and we can only check the equilibrium conditions for a finite subset of these values. We will, however, consider a range of six orders of magnitude in either direction from $\rho = 1$. As we show below, our results clearly converge as $\rho \to 0$ or $\rho \to \infty$, thus providing a high degree of confidence that no relevant case was left unexamined.

We considered the following values for the parameter L: $L = 2, \ldots, 10$. Once a value for L is fixed, our procedure involves three steps.

1. Setting $\gamma = 1$, we performed a coarse search for "problematic cases." This search covered all $d = -L+1, \ldots, L-1$ and $1,000$ different values for ρ on either side of $\rho = 1$. These values were determined so that $\ln(\rho)$ was spread out uniformly between $-\ln(10^6)$ and $+\ln(10^6)$. For each of these (d, ρ)-combinations, we maximized $U_d(s|\rho)$ numerically and checked if the maximum did occur at $s = s_d^*(\rho)$. This condition was *not* satisfied if and only if $d = L-1$ and $\rho < 1$. That is, all cases where the player with fewer remaining resources is one battle away from winning the Tug of War were problematic, and no other cases were.

2. For the problematic cases only, we then performed a second search on a finer grid. This finer search covered $10,000$ different values for ρ, determined so that $\ln(\rho)$ was spread out uniformly between $-\ln(10^6)$ and 0. For each of these cases, we varied γ in increments of 0.001 and determined the highest value for γ such that $U_{L-1}(s|\rho)$ was maximized at $s = s_{L-1}^*(\rho)$. The resulting value, $\overline{\gamma}_L(\rho)$, is shown in Figure 5 and is always between 0.5 and 0.999. Thus, all problematic cases considered turned into well-behaved cases once $\gamma \leq \overline{\gamma}_L(\rho)$.

3. Returning to our coarse grid from step 1 again, we verified that the remaining "non-problematic" cases (i.e., $d < L-1$ and/or $\rho \geq 1$) remained non-problematic when γ was lowered to $0.5, 0.49, \ldots, 0.01$. That is, we checked that $U_d(s|\rho)$ was still maximized at $s_d^*(\rho)$ for these lower values of γ, and determined that this was, indeed, the case.

Note that, in Step 2 of the procedure, once $\rho < 10^{-4}$ the bound $\overline{\gamma}_L(\rho)$ has essentially converged to 0.5 for all $L = 2, \ldots, 10$. Thus, it is implausible that $\overline{\gamma}_L(\rho)$ should fall below

Figure 5: Upper bound for γ such that $U_{L-1}(s|\rho)$ has a global maximum at $s^*_{L-1}(\rho)$.

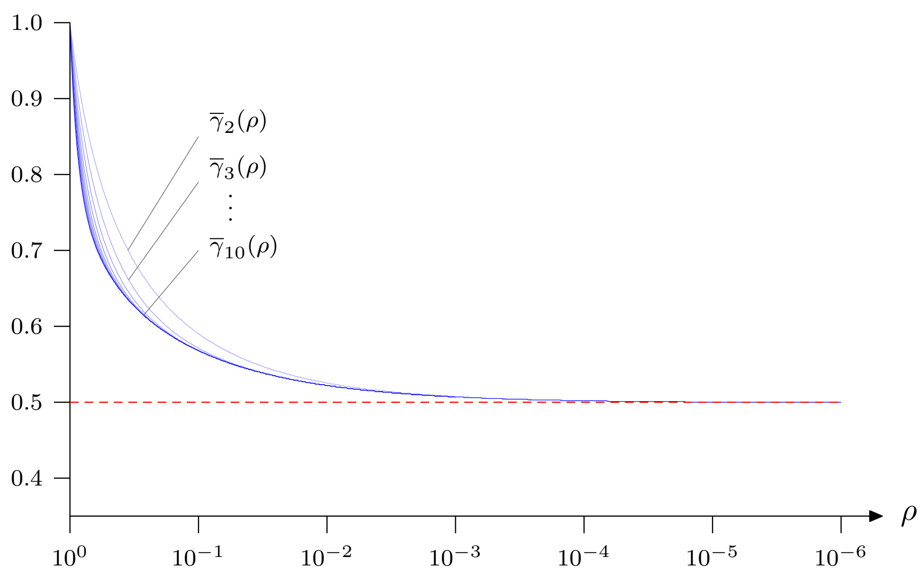

0.5 for values of ρ below 10^{-6}.[16] Therefore, we conclude that balanced strategies given in (20) constitute a subgame perfect equilibrium of the Tug of War game whenever $\gamma \leq 0.5$, but not otherwise.[17]

6 Effort Dynamics in Equilibrium

As we have shown above, in the subgame perfect equilibrium players invest, at every stage where player A is d victories ahead of player B, the fraction $s^*_d(a/b)$ of their remaining resources a and b into the current battle, given by (20). Because $s^*_d(\cdot)$ does not depend on t, the subgame perfect equilibrium is stationary.[18] Because $s^*_d(\rho) = s^*_{-d}(1/\rho)$, the equilibrium is symmetric, and because $s^*_d(\rho)$ depends on a and b only through $\rho = a/b$, the equilibrium is homogeneous.

[16] We also performed Step 2 of the procedure for $L = 50$ and $L = 100$ and confirmed that $\overline{\gamma}_L(\rho)$ converged to 0.5 from above also in these cases. Moreover, $\overline{\gamma}_{50}(\cdot)$ and $\overline{\gamma}_{100}(\cdot)$ were almost identical to $\overline{\gamma}_{10}(\cdot)$.

[17] Similar thresholds arise in adjustable-budget contests that use the Tullock success function. In the basic one-shot Tullock model (which is the special case of a Tug of War with $L = 1$), a pure strategy equilibrium—characterized by the players' respective first-order conditions—exists if and only if $\gamma \leq 2$ (see Baye, Kovenock, and De Vries 1994). In simultaneous majoritarian contests where players must win a majority of $2L - 1$ battles to win the game, a pure strategy equilibrium exists if and only if γ is below some threshold that decreases in L (see Klumpp and Polborn 2006).

[18] This, of course, is hardly a restrictive property—it only means that the contestants apply the same investment functions at any two states (t, d) and (t', d') for which $d = d'$. Since their remaining effort budgets will generally not be the same at states (t, d) and (t', d'), the actual investments made will differ.

The property property that generates the most "structure" is that of balancedness—that is, both player A and player B invest the same fraction, namely $s_d^*(\rho)$. Thus, the players' relative resources stay constant along the equilibrium path and equal to their relative resources at the start of the game, which in turn implies that the players' relative investments are the same in each battle. Under the assumed contest success function, therefore, the probability with which a player wins or loses each component battle of the Tug of War contest does not change as the contest unfolds.[19] Player A wins each battle with probability

$$f(\bar{a},\bar{b}) \;=\; \frac{\rho^{\gamma}}{1+\rho^{\gamma}},$$

with $\rho = \bar{a}/\bar{b}$. The Blotto Tug of War shares this feature of constant battle win probabilities with some other sequential multi-battle contests—in particular, with the non-Blotto, majoritarian two-player contest (where it is true for a broader class of contest success functions; see Klumpp, Konrad, and Solomon 2019) and the non-Blotto, Tug of War team contests with the all-pay auction contest success function (Häfner 2017).

The balancedness property implies that the stochastic process that governs transition from one state to the next in equilibrium is a simple random walk on $\mathbb{Z}$, with incrementation probability $f(\bar{a},\bar{b})$. The expected remaining duration of the contest when player A is d victories ahead of player B is given by

$$D_d^*(\rho) \;=\; \begin{cases} \dfrac{1+\rho^{\gamma}}{1-\rho^{\gamma}}\left[2L\dfrac{1-\rho^{\gamma(L-d)}}{1-\rho^{2\gamma L}} - L + d\right] & \text{if } \rho \neq 1, \\ (L+d)(L-d) & \text{if } \rho = 1 \end{cases} \tag{22}$$

(see Stern 1975; see also Häfner 2017, Corollary 2). Comparing (20) and (22), we can write the fractional equilibrium investment as

$$s_d^*(\rho) \;=\; \frac{1}{D_{-d}^*(\rho)}.$$

Note that, when the budgets are identical (i.e., $\rho = 1$), we have $D_d^*(1) = D_{-d}^*(1)$, which means that $s_d^*(1)$ is simply the inverse expected remaining duration of the contest. In this symmetric case, the equilibrium can be interpreted as players "stretching" their remaining budgets to last the expected remaining duration of the Tug of War. This behavior is somewhat reminiscent of the equilibrium in the Blotto majoritarian contest, where at every stage players stretch their remaining budgets to last the *maximum* remaining duration of the contest (Klumpp, Konrad, and Solomon 2019). However, this similarity

[19]However, should one player deviate from the balanced equilibrium profile, relative budgets will change following the deviation, and so will the probability with which the player wins each subsequent battle.

Figure 6: Equilibrium efforts in the Blotto Tug of War ($L = 6$, $\gamma = 0.5$).

(a) Equilibrium effort as a percentage of remaining budget

(b) Expected equilibrium effort as a percentage of initial budget

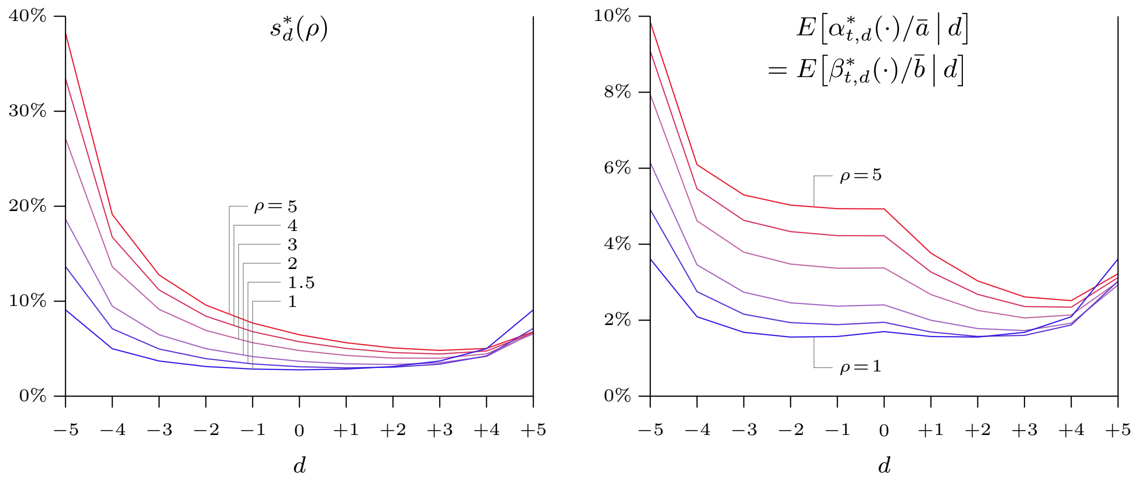

Notes: Panel (a) contains the precise values of $s_d^*(\rho)$, given in (20). Values in panel (b) is based on 10^6 simulated Blotto Tug of War contests for each value of ρ considered.

does not extend to the case where budgets are not identical (i.e., $\rho \neq 1$). In general, what matters is the expected remaining duration of a contest in which the budgets are the same but the lead is *reversed*. That is, if player A is d victories ahead of player B, players compute the expected remaining duration of the Tug of War assuming that B is d victories ahead of A, and then determine their investments by stretching their budgets over this expected duration.

For the case $L = 6$ and $\gamma = 0.5$, Figure 6 (a) plots the relative investments $s_d^*(\rho)$ for six different values of ρ.[20] When the game is entirely symmetric ($\rho = 1$) efforts are between 2.78 percent and 9.09 percent of the players's remaining budgets, and lowest when $d = 0$. Percentage efforts escalate when the game moves closer to one player's line of victory (i.e., $d < 0$ or $d > 0$). This escalation is especially pronounced in cases where player A has a strictly larger budget ($\rho > 1$) and player B comes close to winning the game (i.e., $d < 0$). If one contestant has five times the budget of his opponent and the opponent is one battle away from winning the Tug of War—which, of course, is a relatively unlikely event—the players allocate more than 38% of their remaining budgets to this potentially decisive battle. In contrast, when the player with the larger budget comes closer to winning the game, percentage efforts first decrease (at small positive d) before increasing slightly.

[20]Figure 6 show efforts for $\rho \geq 1$ only, i.e., when player A has the weakly larger starting budget. However, since $s_d^*(\rho) = s_{-d}^*(1/\rho)$, the case where player B has a larger budget can be obtained by replacing ρ with $1/\rho$ and switching the sign of d.

Figure 6 (b) plots the percentage of a player's *initial* budget that is invested into battles where player A is d victories ahead of player B. This is an expectation because, for the same d, a player's absolute effort will depend on the player's remaining budget, which decreases over the course of the game. The game may "cycle" through the same value d multiple times, with the contestants' remaining budgets being lower than before on each occasion. To obtain an expected value of this percentage, we simulated the Blotto Tug of War one million times for each value of ρ considered. Similar to efforts as a percentage of remaining budget, efforts as a percentage of initial budgets escalate in the event the richer contestant falls behind.

Figure 7 plots the expected percentage of a player's initial budget that is invested at stage t of the contest, conditional on that stage being reached. This percentage, too, was obtained by simulating the Blotto Tug of War one million times for each value of ρ considered. Because players have only limited budgets and invest at least some small percentage of their remaining resources into each battle, absolute investments must decline in the long run. However, for a symmetric contest (i.e., $\rho = 1$) this decline is very slow; that is, the effort investment profile over time is relatively flat. The intuition is that each player wins each battle with probability 1/2. Therefore, the contest will remain at relatively symmetric states (i.e., d close to zero) for a relatively long time, and players invest only a small percentage of their budgets in these states. Effort escalation occurs at relatively asymmetric states, which are reached relatively late and, hence, with fewer remaining resources. Effort declines more steeply when the contest is more asymmetric (i.e., $\rho > 1$). In this case, players invest a larger percentage in the more symmetric states,

Figure 7: Equilibrium efforts over time ($L = 6$, $\gamma = 0.5$).

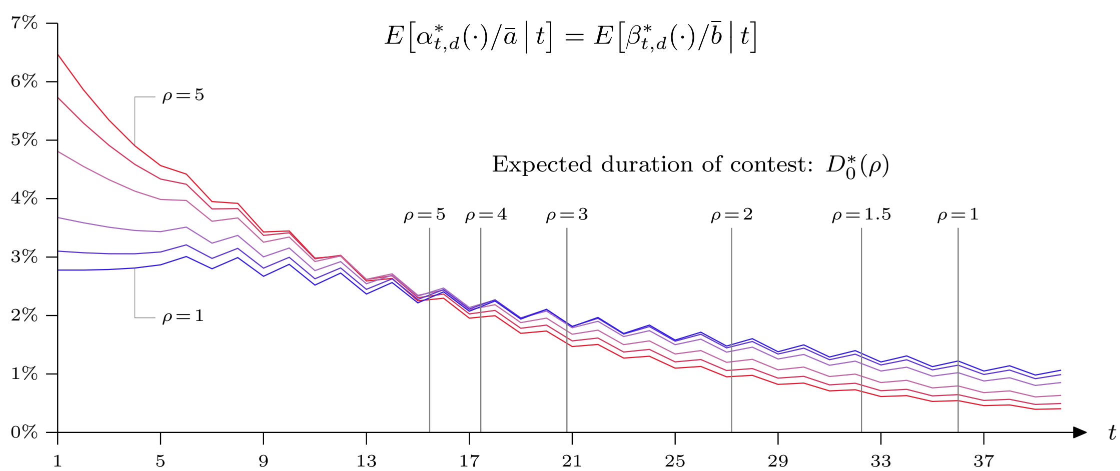

Notes: Values are based on 10^6 simulated Blotto Tug of War contests for each value of ρ considered. Expected durations are computed using (22).

in which the contest is more likely at an early time than at a later time. Moreover, the asymmetric states where the equilibrium prescribes severe effort escalation are those where the richer player has fallen behind, but these states are unlikely to be reached at all when initial budgets are highly asymmetric.[21]

Finally, Figure 7 also shows the expected duration of the Blotto Tug of War, $D_0^*(\rho)$, for each value of ρ considered. More asymmetric contests end more quickly on expectation. When adjusting for these differences in duration, effort decline rates are much more uniform: Effort at the expected end of the Tug of War (rounded to the nearest integer) is between 60 and 65 percent lower than effort in the initial period, in all six cases considered. The percentage of their initial budgets that players have spent, on average, by the end of the Tug of War varies from 56 percent (for $\rho = 5$) to 64 percent (for $\rho = 1$). Thus, despite resources having no residual value outside of the Blotto contest, players leave more than one third of their initial budgets unused on expectation.

7 Conclusion

This is, to our knowledge, the first paper to examine Tug of War contests with the Blotto budget specification. The main complication one must deal with when analyzing this game is that players have finite resource budgets which they must allocate to potentially infinitely many battles. To deal with this complication, we assumed that each battle in the Tug of War is represented by a (sufficiently concave) Tullock contest success function and that players do not discount future payoffs. Under these assumptions, we were able to derive a subgame perfect equilibrium strategy in closed form. The key property of this equilibrium is that the two players invest identical fractions of their remaining resource budgets at each stage. While this fraction depends on the lead one player has over the other, the fact that both players invest the same fraction implies that win probabilities in individual battles of the Tug of War remain constant.

One natural extension of the model is to envision an extended game in which the two players first raise resources at a cost, and then compete with these resources in the Tug of War. Using the two-firm example in the Introduction, the startups could first compete in a capital raising round for funding, and then compete for market share in the Tug of War. Using (11), the *ex ante* probability that player A wins the Blotto Tug of War is

$$u_0(\bar{a}/\bar{b}) \;=\; \frac{(\bar{a}/\bar{b})^{\gamma L}}{1 + (\bar{a}/\bar{b})^{\gamma L}}.$$

[21]The curves in Figure 7 are not monotonically decreasing, but have small "bumps." Recall that, while relative investments $s_d^*(\rho)$ depend on d and not not t, the variables d and t are related: d must be even whenever t is odd, and vice versa. Thus, the lead that one player has over another is restricted to a different set of values in odd periods than it is in even period, and so are the values that $s_d^*(\rho)$ takes.

While we have not proven uniqueness of the equilibrium, the standard properties of constant-sum games imply player A wins the Tug of War with probability $u_0(\bar{a}/\bar{b})$ in *any* equilibrium, whenever the players' starting budgets are $\bar{a}$ and $\bar{b}$. Letting c_i denote the marginal fundraising cost for $i = A, B$, the players choose $\bar{a}$ and $\bar{b}$ in the initial stage of the extended game to maximize

$$2u_0(\bar{a}/\bar{b}) - 1 - c_A\bar{a} \quad \text{and} \quad 1 - 2u_0(\bar{a}/\bar{b}) - c_B\bar{b},$$

respectively (assuming they obtain payoff 1 if they win the Tug of War, and payoff -1 if they lose).

For given γ, $\bar{a} > \bar{b}$ implies $u_0(\bar{a}/\bar{b}) \to 1$, and $\bar{a} < \bar{b}$ implies $u_0(\bar{a}/\bar{b}) \to 0$, as $L \to \infty$. Thus, for large L, even a small resource advantage over the opponent allows a player to win the Tug of War with a probability of close to one. The arguments in Klumpp, Konrad, and Solomon (2019, Section 5) then imply that the fundraising stage of the game is approximately an all-pay auction, in which players adopt mixed strategies in equilibrium. In particular, if $c_B \geq c_A$, player A randomizes $\bar{a}$ uniformly on $[0, 2/c_B]$, while and player B sets $\bar{b} = 0$ with probaility $1 - c_A/c_B$, and randomizes $\bar{b}$ uniformly on $[0, 2/c_B]$ with the remaining probability.

Appendix

Proof of Lemma 2

Let $\mu = a^\gamma/(a^\gamma + b^\gamma)$. The problem of computing u_d is a variant of the gambler's ruin problem, that is, determining the probability that a simple random walk on $\mathbb{Z}$ that begins at d and increments with probability μ (i.e., drift $2\mu - 1$) hits L before it hits $-L$. To do so, let u_d denote this probability. Note that

$$u_d \;=\; \mu u_{d+1} + (1-\mu)u_{d-1} \;\Rightarrow\; \big[u_{d+1} - u_d\big] \;=\; \frac{1-\mu}{\mu}\big[u_d - u_{d-1}\big]. \tag{23}$$

Solving this difference equation with the boundary conditions $u_{-L} = 0$ and $u_L = 1$ gives

$$u_d \;=\; \begin{cases} \dfrac{1 - \left(\frac{1-\mu}{\mu}\right)^{L+d}}{1 - \left(\frac{1-\mu}{\mu}\right)^{2L}} & \text{if } \mu \neq 1/2, \\[2ex] \dfrac{L+d}{2L} & \text{if } \mu = 1/2 \end{cases}$$

$$= \begin{cases} \dfrac{\left(\frac{\mu}{1-\mu}\right)^{L-d} - \left(\frac{\mu}{1-\mu}\right)^{2L}}{1-\left(\frac{\mu}{1-\mu}\right)^{2L}} & \text{if } \mu \neq 1/2, \\ \dfrac{L+d}{2L} & \text{if } \mu = 1/2. \end{cases}$$

Since $\mu/(1-\mu) = (a/b)^\gamma = \rho^\gamma$, the formula (11) follows. □

Derivation of $u_d'(1)$ for Section 5.2

To see that (11) is continuous at $\rho = 1$, take

$$\begin{aligned} \lim_{\rho\to 1,\, \rho\neq 1} u_d(\rho) &= \lim_{\rho\to 1} \frac{\rho^{\gamma(L-d)} - \rho^{2\gamma L}}{1-\rho^{2\gamma L}} \quad (\to \text{“}0/0\text{”}) \\ &= \lim_{\rho\to 1} \frac{\gamma(L-d)\rho^{\gamma(L-d)-1} - 2\gamma L\rho^{2\gamma L-1}}{-2\gamma L\rho^{2\gamma L-1}} = \frac{L+d}{2L} = u_d(1), \end{aligned}$$

where the second equality applies H'ôpital's Rule. If $u_d(\rho)$ is differentiable at $\rho \neq 1$ and $\lim_{\rho\to 1,\, \rho\neq 1} u_d'(\rho)$ exists, then $u_d(\rho)$ is differentiable at ρ, with $u_d'(1) = \lim_{\rho\to 1,\, \rho\neq 1} u_d'(\rho)$.

For $\rho \neq 1$, it will be convenient to express $u_d(\rho)$ as

$$u_d(\rho) = \frac{\rho^{\gamma(L-d)} - 1}{1-\rho^{2\gamma L}} + 1.$$

Then, for $\rho \neq 1$, we have

$$\begin{aligned} u_d'(\rho) &= \frac{\gamma(L-d)\rho^{\gamma(L-d)-1}\left[1-\rho^{2\gamma L}\right] - \left[\rho^{\gamma(L-d)} - 1\right](-2\gamma L\rho^{2\gamma L-1})}{\left[1-\rho^{2\gamma L}\right]^2} \\ &= \gamma\frac{(L-d)\rho^{\gamma(L-d)-1} + (L+d)\rho^{\gamma(3L-d)-1} - 2L\rho^{2\gamma L-1}}{\left[1-\rho^{2\gamma L}\right]^2}. \end{aligned}$$

Therefore,

$$\begin{aligned} \lim_{\rho\to 1,\, \rho\neq 1} u_d'(\rho) &= \lim_{\rho\to 1} \gamma\frac{(L-d)\rho^{\gamma(L-d)-1} + (L+d)\rho^{\gamma(3L-d)-1} - 2L\rho^{2\gamma L-1}}{\left[1-\rho^{2\gamma L}\right]^2} \quad (\to \text{“}0/0\text{”}) \\ &= \lim_{\rho\to 1} \frac{\begin{array}{c}(L-d)(\gamma(L-d)-1)\rho^{\gamma(L-d)-2} \\ +\,(L+d)(\gamma(3L-d)-1)\rho^{\gamma(3L-d)-2} \\ -\,2L(2\gamma L-1)\rho^{2\gamma L-2}\end{array}}{-4L\rho^{2\gamma L-1} + 4L\rho^{4\gamma L-1}} \quad (\to \text{“}0/0\text{”}) \end{aligned}$$

$$= \lim_{\rho\to 1} \frac{\begin{array}{c}(L-d)(\gamma(L-d)-1)(\gamma(L-d)-2)\rho^{\gamma(L-d)-3}\\ +(L+d)(\gamma(3L-d)-1)(\gamma(3L-d)-2)\rho^{\gamma(3L-d)-3}\\ -2L(2\gamma L-1)(2\gamma L-2)\rho^{2\gamma L-3}\end{array}}{-4L(2\gamma L-1)\rho^{2\gamma L-2}+4L(4\gamma L-1)\rho^{4\gamma L-2}}$$

$$= \frac{\gamma^2\left[2L^3-2Ld^2\right]}{8\gamma L^2} = \gamma\frac{(L+d)(L-d)}{4L}.$$

where the second and third equalities apply H'ôpital's Rule. It follows that $u_d(\rho)$ is differentiable at $\rho = 1$, with $u_d'(\rho) = \gamma(L+d)(L-d)/(4L)$.

References

[1] Agastya, Murali and R. Preston McAfee. 2006. "Continuing Wars of Attrition." Working paper.

[2] Baye, Michael, Dan Kovenock, and Casper De Vries. 1994. "The Solution to the Tullock Rent-Seeking Game when $R > 2$: Mixed Strategy Equilibria and Mean Dissipation Rates." *Public Choice* 81, 363–380.

[3] Borel, Émile (1921). "La theorie du jeu les equations integrales a noyau symetrique." *Comptes Rendus del Academie* 173, 1304–1308. English translation by Savage, Leonard (1953), "The Theory of Play and Integral Equations With Skew Symmetric Kernels," *Econometrica* 21, 97–100.

[4] Budd, Christopher, Christopher Harris, and John Vickers. 1993. "A Model of the Evolution of Duopoly: Does the Asymmetry between Firms Tend to Increase or Decrease?" *Review of Economic Studies* 60 (3), 543–573.

[5] Feller, William. 1970. *An Introduction to Probability Theory and Its Applications.* New York City, NY: Wiley.

[6] Fu, Qiang, Jingfeng Lu, and Yue Pan. 2015. "Team Contests with Multiple Pairwise Battles." *American Economic Review* 105 (7), 2120–2140.

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

[8] Gelder, Alan. 2014. "From Custer to Thermopylae: Last Stand Behavior in Multi-Stage Contests." *Games and Economic Behavior* 87, 442–466.

[9] Häfner, Samuel. 2017. "A Tug of War Team Contest." *Games and Economic Behavior* 104, 256–274.

[10] Häfner, Samuel. 2020. "Eternal Peace in the Tug-of-War?" *Economic Theory*, forthcoming.

[11] Harris, Christopher and John Vickers. 1987. "Racing with Uncertainty." *Review of Economic Studies* 54 (1) 1–21.

[12] Klumpp, Tilman and Mattias Polborn. 2006. "Primaries and the New Hamshire Effect." *Journal of Public Economics* 90, 1073–1114.

[13] Klumpp, Tilman, Kai Konrad, and Adam Solomon. 2019. "The Dynamics of Majoritarian Blotto Games." *Games and Economic Behavior* 117, 402–419.

[14] Konrad, Kai. 2009. *Strategy and Dynamics in Contests*. Oxford University Press.

[15] Konrad, Kai and Dan Kovenock. 2005. "Equilibrium and Efficiency in the Tug of War." Working paper.

[16] Konrad, Kai and Dan Kovenock. 2009. "Multi-Battle Contests." *Games and Economic Behavior* 66, 256–274.

[17] Kovenock, Dan and Brian Roberson. 2010. "Conflicts With Multiple Battlefields." CESifo working paper: Empirical and Theoretical Methods, No. 3165.

[18] Kvasov, Dmitriy. 2007. "Contests With Limited Resources." *Journal of Economic Theory* 136 (1), 738–748.

[19] Malueg, David and Andrew Yates. 2010. "Testing Contest Theory: Evidence from Best-of-Three Tennis Matches." *Review of Economics and Statistics* 92 (3), 689–692.

[20] Roberson, Brian. 2006. "The Colonel Blotto Game." *Economic Theory* 29 (1), 1–24.

[21] Roberson, Brian and Dmitriy Kvasov. 2012. "The Non-Constant-Sum Colonel Blotto Game." *Economic Theory* 51, 397–433.

[22] Sela, Avner. 2011. "Best-of-Three All Pay Auctions." *Economics Letters* 112, 67–70.

[23] Stern, Frederick. 1975. "Conditional Expectation of the Duration in the Classical Ruin Problem." *Mathematics Magazine* 48 (4), 200–203.

[24] Tullock, Gordon. 1980. "Efficient Rent-Seeking." In: James Buchanan, Robert Tollison, and Gordon Tullock (eds.), *Towards a Theory of the Rent-Seeking Society*, 97–112. College Station, TX: Texas A&M University Press.