The Pennsylvania State University

The Graduate School

Department of Political Science

THE COMPLEXITY OF CONFLICT MANAGEMENT

A Dissertation in

Political Science

by

Anna O. Pechenkina

© 2013 Anna O. Pechenkina

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

December 2013

The dissertation of Anna O. Pechenkina was read and approved* by the following:

D. Scott Bennett
Distinguished Professor of Political Science
Head of the Department of Political Science
Dissertation Adviser
Chair of Committee

Douglas Lemke
Associate Professor of Political Science

Philip A. Schrodt
Professor of Political Science

John Yen
University Professor of Information Science and Technology
Affiliated Professor of Computer Science Engineering

*Signatures on file in the Graduate School.

Abstract

The literatures on conflict management and intervention have developed in isolation, rather than informing each other. This project initiates the bridging of the two lines of research by representing mediation and external support in the same agent-based computational model. Two questions motivate my dissertation. First, I explore the effect of external support on war duration by simulating the underlying bargaining process in the conflictual and peaceful interactions of a multi-player international system. The empirical research on external support in conflict has had inconsistent results.

I follow previous research in assuming that third parties choose to reinforce either side when it can make a difference in the outcome, which means support is never given to those belligerents that can win regardless. Given that third parties in my model observe the behavior of supporters for the opposing side, I also derive that support is more likely to be given when one's protege's opponent receives support. This in turn implies that support on average makes wars longer because it is meant to alter the outcome of the conflict. In departure from previous research, my model also accounts for the presence of power mediators who may impose prohibitive costs on the belligerents and may make the provided support irrelevant. Finally, my model also yields that support may alter the calculus of belligerents during the conflict of interests stage, preventing the conflict from escalating to violence in the first place. Therefore, the effect of support on conflict duration is bimodal: preventing the onset and prolonging those armed conflicts that have started.[1]

[1] Preventing the onset may be viewed as a decrease in duration from a positive number of armed conflict months to zero months.

Second, I focus on the effect of power mediation (i.e., a mediation style that exerts pressure on the belligerents) on peace duration. Power mediation is the most effective mediation style at achieving ceasefires. This finding may have important policy implications: the more pressure mediators put on the belligerents the more likely they are to precipitate peace. However, power mediation may also reduce the duration of ceasefire that follows, which implies that third parties may sacrifice long-term peace stability for the sake of short-term benefits of an unstable ceasefire. I demonstrate that the actual effect of power mediation on ceasefire survival is more nuanced and should be analyzed in concert with the actions of support-providing third parties. First, if powerful mediators stop inflating the costs of fighting the conflict of interest and armed conflict may indeed recur in my model, as previous scholarship argues. However, this dynamic takes place in those cases that receive no external support. When reinforcement is large enough to eliminate uncertainty over who is going to win the war, my model yields that a peaceful transfer of resources to the dissatisfied belligerent will take place. In contrast, in those cases when reinforcements are only large enough to recreate conflicts of interests but do not remove uncertainty, an initially fragile ceasefire may become even more prone to conflict recurrence than previously argued. Therefore, the effect of power mediation on ceasefire duration depends heavily on the behavior of support providing third-parties.

The flexibility of computer programming to accommodate a finite number of heterogeneous actors and my focus on individual-level decisions make agent-based modeling a suitable tool for hypotheses development in this project. I build a benchmark agent-based model of intra-war and post-war bargaining in the world without third parties in chapter 2 and an extended model with mediating and support-providing third parties in chapter 4. One of the disadvantages of computational models is that parameter spaces are often prohibitively large to generate and analyze exhaustively. This results in scholars often ignoring the boundary conditions for a model's parameters and non-systematically derived conclusions. I introduce a systematic way to interpret agent-based models by interfacing a genetic algorithm search program with simulation models, conceptualizing numerical output from a simulation model as a fitness function of the genetic algorithm. Genetic algorithms do not need to create an exhaustive parameter space to find global solutions, and therefore users are able to explore their models more quickly, yet systematically. The software program *SimGA* is presented in chapter 3 and verified by recreating the logic underlying the benchmark model of chapter 2. I also use this program to analyze the extended model in chapter 4.

Finally, I use observational data to test empirically some of the extended model's implications in chapter 5. I test for effect of external support on war duration against a sample of civil wars and find support for my model's implications. With respect to external support shaping peace duration, in the sample of interstate ceasefires, external support precipitates conflict when it recreates uncertainty and stabilizes peace when it eliminates uncertainty.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Overview of the Project

I represent mediation and intervention in the same model to bridge the two literatures that have developed largely in isolation from each other. In press, the policy implications of the decisions to reinforce and/or pressure for peace are often viewed as evident,[1] despite the fact that academic research has had inconclusive results about how the strategies of conflict management affect war and peace processes. Two major questions motivate this project. First, what is the effect of external support on war duration? Second, what is the effect of power mediation (i.e., a mediation style that exerts pressure on the belligerents) on peace duration?

My first question is impelled by the contradictory findings in the literature on external support in civil wars: some authors argue that external support extends conflict duration (Regan 2000, 2002; Regan and Aydin 2006), other studies find that external support shortens war (Balch-Lindsay, Enterline and Joyce 2008; Collier, Hoeffler and Soderbom 2004). Gent (2008) argues that third parties select their recipients carefully, which in turn affects how support influences war duration until a certain outcome (in the context of civil wars, Gent argues, support helps rebels win, but has no effect on government's victory).

[1] E.g., see Entous (June 15, 2013), Zakaria (June 20, 2013) or Headlee (June 19, 2013) for the discussion of potential implications of the US's diplomacy and weapons shipments in Syrian civil war in addition to other state and non-state actors' involvement in that conflict. All three examples imply that the suggested consequences of third-party actions are very likely to take place.

In the interstate war literature, most researchers focus only on the questions of why sometimes third parties put boots on the ground or "join" wars, while staying out in other cases.[2][3] While significant transfers of weapons and equipment are commonplace in interstate wars,[4] there has been a lack of quantitative research on the effect of such transfers on interstate war duration. Since I assume that the underlying logic of conflict is analogous in both intra- and inter-state wars, I expect my model's predictions to apply to both "types" of conflict as long as I account for the characteristics that are more often observed in some conflicts rather than others.[5]

I build on Gent's (2008) work by incorporating the supply side decisions in my agent-based model. Assuming that third parties provide support to achieve a certain distributional outcome, it is never given to those belligerents who can arrive at the sought outcome regardless. Given that third parties in my model may reinforce both sides of the warring dyad, only the net shift in capabilities matters. Support that eliminates uncertainty shortens conflict, while support that (re-)creates uncertainty prolongs it. Finally, my model also yields that support may alter the calculus of belligerents during the conflict of interests stage, preventing the conflict from escalating to violence in the first place.

My second question is animated by the tension in the literature on peace duration. Multiple studies find that power mediation is the most effective mediation style at achieving ceasefires (Wilkenfeld et al. 2003; Schrodt and Gerner 2004; Beardsley et al. 2006; Bercovitch

[2] See Joyce, Ghosn and Bayer (2013) for a detailed literature review on this question.

[3] The effect of multiple states participating in war as opposed to a two-side war duration has been shown to reduce interstate wars (Bennett and Stam 1996). Participation of multiple states in an interstate war should not be viewed as 100% analogous to sending troops to support one of the sides in a civil war, as some of the time multiple states are original belligerents as opposed to joiners.

[4] E.g., consider US supply of weapons and equipment to Israel during the Yom Kippur War in 1973, or the USSR's weapons supply to the Arab states on the opposing side.

[5] See subsection 1.2 for a more detailed explanation of why I consider the logic of intra- and inter-state wars to be analogous.

and Gartner 2009; Beardsley 2011; Quinn et al. 2013) This conclusion may have important policy implications: the more pressure mediators exert on the belligerents the more likely they are to precipitate peace. However, Werner and Yuen (2005) and Beardsley (2011, 2008) find that power mediation reduces the duration of ceasefire that follows. This finding implies that third parties may sacrifice long-term peace stability for the sake of short-term benefits of an unstable ceasefire.

I demonstrate that the actual effect of power mediation on ceasefire survival is more nuanced and should be analyzed in concert with the actions of support-providing third parties. First, if powerful mediators stop inflating the costs of fighting the conflict of interest and armed conflict may indeed recur in my model, as Werner and Yuen (2005) and Beardsley (2011, 2008) argue. However, this dynamic takes place in those cases that receive no external support. When reinforcement is large enough to eliminate uncertainty over who is going to win the war, my model yields that a peaceful transfer of resources to the dissatisfied belligerent will take place. In contrast, in those cases when reinforcements are only large enough to recreate conflicts of interests but do not remove uncertainty, an initially fragile ceasefire may become even more prone to conflict recurrence than previously argued. Therefore, the effect of power mediation on ceasefire duration depends heavily on the behavior of support providing third-parties who may stabilize or undermine ceasefires through external support.

Prior research has explored multiple third-party behaviors, including facilitation of negotiations, manipulation of negotiations, economic intervention, and military intervention. Broadly speaking, all third-party activity may be grouped by its immediate intent: seeking the end of violence vs. seeking a certain distributional outcome. For instance, facilitation of negotiations or pressure for peace normally seek to establish a ceasefire (in its immediate intent). In contrast, providing economic or military support to one of the sides seeks to help the recipient achieve a more favorable settlement than what the recipient would have achieved without such support.[6] With the exception of Regan and Aydin (2006), the studies that focus on mediation and those that explore external support have evolved separately.

The two broad categories of third-party behavior modeled here include pressuring the belligerents to cease fire and providing reinforcements to one of the sides in conflict. I group third-party actions based on their immediate intent instead of considering their long-term goals. In some cases, a third party may seek power mediation, because s/he has both the immediate intent and the long-term goal of peace. In other instances, a third party may use power mediation to stop the fighting, and be mostly concerned with establishing a ceasefire that favors his/her protégé, i.e., a ceasefire that contributes to that third party's geopolitical goals. Similarly, when it comes to external support, Regan (2010, 458) notes that "whether [the support-providing interveners] are to achieve geopolitical objectives or to advance humanitarian considerations such as bringing peace to the country in some cases" has been debated for a long time. Since it is impossible for me to know the "true" goals of third parties, I only use the observable behaviors – power mediation and support – as the basis for the third-party actions in my model.

[6] Regan and Aydin (2006) offer a similar grouping of third-party behaviors, which they label as (1) [third parties] that attempt to influence the structure of the relationship among combatants and (2) [third parties] that attempt to manipulate the information that these actors hold. Regan and Aydin's first group closely mirrors my description of the third parties that seek a certain distributional outcome. Reinforcing one of the sides in the warring dyad shifts the distribution of power in the dyad, and as a result may lead to a desired distribution of benefits as a result of war. Regan and Aydin's latter group describes the third parties that attempt mediation. Their description of mediation as attempts to manipulate information that the belligerents hold hinges on the assumption that mediators have information that is unknown to the belligerents (Kydd 2003). Current research challenges this assumption. In addition, multiple empirical studies show that manipulative mediation (i.e. the use of threats) is the most successful at achieving the end of violence (Wilkenfeld et al. 2003; Schrodt and Gerner 2004; Beardsley et al. 2006; Bercovitch and Gartner 2009; Beardsley 2011; Quinn et al. 2013).

To emphasize the importance of third-party behavior, chapter 2 introduces a model of conflictual and peaceful interactions within dyads excluding third-party actions (also referred to as a "benchmark" or "null" model). In chapter 4, the same model is extended by adding third-party behavior. I analyze the model in two stages for three main reasons. First, the comparison between the models allows me to demonstrate that the fundamentals of conflict and peaceful interactions change profoundly once we account for third party actions. The conflicts that could potentially experience power mediation, last longer in the null model than they do in the extended model. In contrast, the belligerents who have strong outside supporters that would potentially want to reinforce them would fight on average longer wars in the extended model (assuming that on average the losing side of the dyad is the recipient of support). In the null model, the sides settle only when they solve the underlying conflict of interests without being pressured into a ceasefire that does not reflect the true distribution of power in the dyad. This means that ceasefires are stable in the null model, as no random shifts in capabilities occur to upset the established ceasefires. In the extended model, ceasefires are sometimes imposed by power mediators, which creates further recurrence of conflicts of interests when the mediation's effect wanes. Furthermore, provisions of external support and power mediation to other dyads may upset the established settlements in those dyads, for which third parties play the roles of belligerents. Thus, third-party behavior serves as random shocks to capabilities and guarantees conflict recurrence throughout the existence of the system.

Second, given that the major relationships in the null model can be derived analytically, implementing the null model as a computer simulation allows me to verify the computer code. In other words, by demonstrating that the graphical and numerical output from the simulation model implies the same conclusions as the analytical solution, I ensure that the computer code is written correctly.

Third, I verify the effectiveness of the genetic automated search program introduced in chapter 3 to locate global optima by exploring the parameter space of the null model with a genetic algorithm. The proposed software package provides communication between a genetic algorithm search program and simulation models, conceptualizing the output from a simulation model as a fitness function of the genetic algorithm. Since I first create an exhaustive parameter space in section 2.3.3, I am able to compare the results and the costs of interpreting a computational model with the help of a genetic automated search on the one hand and by generating and analyzing an exhaustive parameter space on the other hand. I demonstrate that heuristics-based optimization is an effective and efficient way of interpreting agent-based models.

The project concludes with quantitative tests of some of the extended model's implications against observational data in chapter 5. The extended agent-based model yields four testable hypotheses:

1. Net reinforcements toward greater imbalance shorten war and net reinforcements toward greater parity prolong war.

2. Net reinforcements toward greater imbalance prolong peace and net reinforcements toward greater parity destabilize peace.

3. Net reinforcements toward greater imbalance after an imposed ceasefire decrease the likelihood of armed conflict and increase the likelihood of peaceful renegotiation of the status quo.

4. Net reinforcements in the direction of greater parity after an imposed ceasefire increase an a priori high likelihood of ceasefire failure.

Using a sample of civil conflicts, I test whether net reinforcements in the direction of imbalance shorten conflict, while net reinforcements toward parity prolong it (hypothesis 1). I demonstrate support for the agent-based model's hypothesis through my analysis of reinforcement based on the power structure of the civil war dyad as opposed to the conventional identity-based codings of support (e.g., support to the opposition or government side). Second, I test whether relative reinforcements after a ceasefire is signed have a pacifying or destabilizing effect on ceasefire survival (hypothesis 2). Using a sample of ceasefires established after interstate wars of the 20th century, I find that conflicts of a high level of violence are less likely to recur when shifts toward imbalance happen. Also, when reinforcement creates parity, peace lasts longer.

1.2 Modeling Inter- and Intra-state Conflicts under the Same Framework

Despite multiple authors describing the division between international and civil conflict as misleading,[7] the scientific study of war has been largely separated into "types" of conflict depending on whether the participants are recognized or unrecognized governments in the international system.[8]

[7] See a brief summary of arguments about why division between inter- and intra-state conflicts is arbitrary Cunningham and Lemke (2013, pp. 4-6). In his 2004 paper, Sambanis cautioned that examining a set of civil wars separately from other types of war may lead to biased estimates.

[8] It is important to note that other scholars have argued that application of theoretical arguments about interstate war to a sample of civil wars may not be straightforward. E.g., Regan (2009) makes a case that territorial disputes are central to both inter- and intra-state wars, however territorial civil wars are not fought over territory itself, but rather territory interacts with a concentrated homogenous ethnic group and therefore territory most often is the means to regional independence. Still, a war for regional independence by a homogenous ethnic group would not be waged had the rebelling group benefited from that territory equally or fairly or without "grievance" as Regan himself points out. Similarly, interstate wars over territory are fought to gain more utility from that territory than what was derived before the war (whether the utility comes from resources or from reputation). Therefore, I believe that it

Cunningham and Lemke (2013) review the empirical reasons cited in the literature in support of dividing the conflicts into "types" and establish that civil wars are more frequent, more likely to recur, longer, bloodier, have participants unequal in power, take place on the territory of a single state, have more civilian victims, have higher stakes, and are less likely to end in a negotiated settlement. While the authors acknowledge that these differences may very well exist, they argue and demonstrate empirically that dividing a set of international wars along the same empirical differences eliminates key findings in the literature on the onset of international war (e.g., democratic peace's pacifying effect or parity's increasing effect on war onset).[9] This is an important conclusion, as it points to the possibility of conflict patterns that have never been found only due to the lack of analysis of combined sets of conflicts.

Furthermore, the authors estimate the impact of a conflict being a civil war on the probability of that conflict ending in a peace agreement and find that among conflicts with at least 1,000 battle deaths civil wars are in fact more likely to end in a signed peace deal than interstate wars. This finding would be impossible to establish without combining all conflicts in the same data set.

In sum, Cunningham and Lemke (2013) make an important contribution to the quantitative study of war by demonstrating empirically that divisions of conflicts into intra- and inter-state "types" is arbitrary and potentially inhibits the discovery of important patterns. In this project, I assume that the underlying logic of intra- and post-war bargaining is the

is possible to argue that the difference between territorial issues in civil and international wars are less pronounced than Regan (2009) represents.

[9] It is important to note that there have been important challenges to democratic peace: see Bennett (2006) (on the existence of regime-similarity peace as opposed to democratic peace) and Gartzke (2007) (on the existence of capitalist peace as opposed to democratic peace).

same regardless of whether the belligerents are recognized by the international community. Furthermore, I assume that third parties' decision making about imposing costs for fighting or reinforcing either side of the conflict is analogous for conflicts with established governments and those that are fighting to become such.

1.3 The Use of Computational Modeling

Agent-based modeling is a computational tool that allows researchers to build more realistic worlds than game-theoretic models, with the tradeoff being the lack of formal analytical solutions. Instead, agent-based models (ABMs) are "solved" through numerical experiments for a set of cases. The final result is a potentially more realistic - yet formal - study of a social system (Galán et al. 2009). The analyst defines some characteristics and a range of behaviors for the model's actors, commonly referred to as agents. The researcher can experiment with these artificial worlds through creating and modifying the computer programs (changing the process, the rules by which the agents make decisions, and whether they learn from each other), therefore treating the worlds as computational experiments (Benoit 2001; Cederman 2001). By programming individual decision-making processes, the researcher can observe the "growth" of the population-level dynamics. ABMs may help uncover the mapping of micro-level motives to aggregate outcomes,[10] or help study emergent properties of the system, which are not necessarily traceable to the individual agents that comprise the system (Cederman 2001).

[10] As Schelling (1978) points out in his famous segregation model, focusing solely on the micro-level behavior does not provide insight into the macro-level emergence.

The flexibility of computer programming to accommodate a finite number of heterogeneous actors and the focus on individual-level decisions make agent-based modeling a suitable tool for hypotheses development in this project. I am interested in understanding the emergent outcomes within the international system that has multiple heterogeneous agents instead of a few or infinite number of agents as would be possible to handle with a game-theoretic model (Miller and Page 2007). Furthermore, in the extended version of the model, I include the behavior of multiple third parties managing multiple conflicts that simultaneously take place in the international system. Each third party makes a decision whether to pressure the belligerents for peace or support one of the sides in an ongoing conflict. A model with so many actors would be intractable or very difficult to solve analytically. The ABM approach also allows me to uncover the implications of management attempts by a cluster of actors as opposed to conflict management by one actor at a time.

Chapter 2

A Benchmark Agent-Based Model of Intra- and Post-War Bargaining

In this chapter, I present the baseline model of conflict and post-conflict bargaining within dyads of recognized and unrecognized governments. The null model assumes away any influence external actors may have on conflict processes. The model yields that war is an effective instrument for achieving global peace when no changes to capabilities occur. Section 2.1 lays out the bargaining relationships within conflictual dyads analytically. I rely on the insights from the bargaining literature to model the process of violent conflict, subsequent spells of peace, and potential conflict recurrence. Section 2.2 presents the details of the computer code that implements the aforementioned analytical model as an agent-based computational model of conflict management. Finally, section 2.3 presents the details of a single run of the model and quantitative analysis of the batch of model runs. I compare the conclusions derived analytically to the numerical output of the simulation and therefore verify that the computer code is implemented correctly.

2.1 Theoretical Foundations of the Null Model

In this section, I lay out a simple bargaining model of conflict initiation, duration, and recurrence. The agents in my model represent recognized and/or unrecognized governments that sometimes have conflicts of interests that may develop into armed conflicts due to uncertainty about how much relative power belligerents have. I incorporate the idea from

Powell (1996, 1999, 2004) that the probability of war onset is proportional to the gap between the distribution of power and the distribution of benefits. The distribution of power is represented by relative military capabilities and is reflective of the sides' abilities to win a war. The distribution of benefits is represented by relative allocation of the disputed good. Once the benefits that one side derives from the disputed good do not sufficiently reflect one's power, war may become attractive as one may gain more benefits from fighting a war. Fearon (1995) posits an important question of why states pay the cost of fighting a war instead of dividing the good through negotiation and avoiding the cost of war. The answer, he argues, is in the uncertainty about relative capabilities.[1] If states are uncertain about how strong their opponent is, they may overestimate their chances of victory and thus rationally choose to fight. Powell (1996) adds to Fearon's (1995) insights that information problems do not lead to war independent from the current distribution of benefits. It is only when one can expect to gain from war, uncertainty may lead to armed conflict. Powell's argument has received strong empirical support in Reed et al. (2008).

The model is structured by phases: war and peace. A dyad enters the *war phase* if there exists a conflict of interests, such that one side has an incentive to challenge the opponent who controls a greater share of the prize, than what is proportional to that side's military might. I model the war phase as a "learning equilibrium" (Powell 2004, 348-349), which assumes that sides fight as long as their expectations about winning the next battle diverge (Gartzke 1999).[2] In this project, I assume that states fight when the defender is

[1] Fearon (1995) distinguishes between information-, and commitment-driven wars. Powell (2006) has shown that Fearon's (1995) third type of war, indivisibility-driven war, is a subset of the commitment-driven wars.

[2] As opposed to wars, in which states have to fight until the full victory or the physical destruction of the source of the advantage, which is the basis for the commitment problem (Gartzke 1999).

uncertain over the challenger's capabilities.[3] The challenger side has a private capability component in addition to observable capabilities.

All states in the system are paired into non-directional dyads, each of which has side A and side B.[4] Suppose π denotes the distribution of benefits in the dyad and reflects how much of the disputed good A controls. Assuming that the good is finite (total good equals 1), B controls $(1 - \pi)$. A gap between the distribution of power and the distribution of benefits arises when the amount of benefits controlled by A exceeds her power.[5] I return to a more specific description of what constitutes a conflict of interests once I discuss the sides' probabilities of winning militarily.

War phase: In the model, during a war, side A is granted with the power to propose the terms of the settlement x_t. Let $x_t \in [0, 1]$ denote the proposal made at time t. Side B has the power to accept or reject the terms. Terms or an "offer" is A's best assessment of how much of the prize A can afford to keep. Therefore, smaller offers are more generous as they imply that weak A intends to keep less, and vice versa, greater offers are more demanding as they imply that A plans to keep more of the prize. If B accepts A's offer, the sides do not fight a battle and settle if A is certain about B's type.[6] B's rejection sends a signal that B is more powerful than A thinks and that B is willing to fight at least another battle to get a better settlement. A then needs to keep less of the prize in order to satisfy stronger B. If the offer

---

[3] Another plausible modeling choice is to assume uncertainty over resolve (Fearon 1995).

[4] Following a weak convention in the literature, I refer to A as "she" and to B as "he".

[5] As explained further below in this section B is assumed to be a challenger and A is assumed to be the defender of the prize.

[6] There exists a scenario of B accepting an offer that may satisfy both a weak and a moderately strong B, thus conflict of interests (although not an armed conflict) persists in the following round. In other words, settlement implies B's acceptance of A's terms and A having no uncertainty over B's type.

is separating between B's types,[7] then A's beliefs are updated accordingly. After an offer is rejected, the sides fight a battle. If the previous round's offer was not separating, A's beliefs about whether B is strong or moderately strong are updated based on the battle outcome. In the following round, A makes her proposal more generous in attempt to match stronger B's expectations. When A is uncertain between two types of B the offer is separating by definition, therefore the second battle is always the final one in this model; this is a function of my choice to assume that B may be of three types.[8] If the previous round's offer was separating, then A's beliefs are updated before the sides fight a battle and the battle is merely a mechanism to distribute the good.

By accepting the terms, B signals that he is either as powerful as the offer implies or that B could potentially be weaker than A initially thought. In some cases, it is possible for B to accept the terms and for A to continue revisiting the settlement by making another offer if A believes that B has revealed that he is weak and A thinks she could gain more of the prize.[9] Thus if B accepts a non-separating offer, the conflict of interest persists and A makes another offer in the following round, which risks war.

In sum, A's offers about what proportion of the good A would like to keep are manifestations of A's beliefs about B's strength, while B's decisions about acceptance or rejection of the terms are signals about B's actual strength. Several variations of this "learning through

[7] A separating offer is an offer, the rejection of which reveals B's type fully. E.g., a separating offer is such that is only rejected by a strong B, therefore if B rejects, the rejection solves uncertainty and A's beliefs about the probability that B is strong are updated accordingly (both A's belief that B is weak and A's belief that B is moderately strong are set to 0. A non-separating offer is an offer, the rejection of which narrows the possible number of B's types. E.g., an offer that is only acceptable to a weak B, when rejected reveals that B is either moderately strong or strong, thus narrowing the number of possible types from three types to two possible types. A's belief that B is weak is then updated to equal 0.

[8] This is an arbitrary modeling choice. Assuming four types of B would allow for conflicts of a maximum of three battles, etc. The goal of this model is not to describe reality perfectly, but instead to convey the logic of why sometimes uncertainty-driven wars are short, while in other cases they take a long time to end.

[9] E.g., by accepting an offer that is acceptable to both a weak B and a moderately strong B, B signals he is not strong, therefore A's beliefs about ϕ_w and ϕ_m are updated accordingly.

fighting" concept have been modeled in the bargaining game-theoretic literature (Wagner 2000; Filson and Werner 2002; Slantchev 2003; Smith and Stam 2004).

A and B are characterized by their capabilities. The probability of winning a battle is assumed to be a function of capabilities. For instance, A's probability of winning a battle in round t is p_t equals the proportion of the capabilities in the dyad that A controls, $p_t = \dfrac{m_A}{m_A + m_B}$, where $m_A > 0$ denotes A's capabilities and m_B are B's. B's probability of winning a battle equals $1 - p_t$. I assume that there exists uncertainty over what proportion of B's total capabilities B is concealing.[10]

Since A is the offer-making side and B is granted with the power to accept or reject offers in this setup, I focus on B's utility. B compares his value of accepting an offer to the known payoff from fighting a battle, which means that B is willing to keep at least $1-p_t-c_B$ portion of the prize, where c_B denotes the cost that B endures from fighting a battle. Since A's probability of winning is p_t, B will win a battle with probability $1 - p_t$, therefore, B prefers anything that is at least as good (for B) as $1 - p_t - c_B$.

Three types of B: As mentioned above, I assume that A is uncertain about the component of B's capabilities that is private. I assume that A may have three guesses about B's private capabilities. Conceptually, these three guesses, modeled as three values, represent a lower, a medium, and an upper bound of B's private capabilities.[11] Suppose that B's total capabilities equal $m_{Bt} + b$, where $m_{Bt} > 0$ is a publicly observable component of B's

[10] One may think of this scenario as if B may attempt to conceal some of his capabilities, or A may be wrong in her estimates as not all of B's capabilities are perfectly observable. Empirical examples of this dynamic include Iran's attempts to conceal its nuclear program and the attempts by other states to attain intelligence about the state of Iran's program, which by definition may not be fully accurate.

[11] In the model, the three values are drawn from a uniform or normal distribution. To ensure that the choice of the distribution is not affecting the results, I use different distributions across simulations.

capabilities (which can change over time as a function of third party reinforcement in the extended model), and $b > 0$ is the private component of B's capabilities, which is constant over time. Let $b = \underline{b}$ with probability ϕ_w, $b = \hat{b}$ with probability ϕ_m, and $b = \bar{b}$ with probability $1 - \phi_w - \phi_m$, where ϕ_w and ϕ_m are exogenously given beliefs of A about the value of b and $\underline{b} < \hat{b} < \bar{b}$. Suppose that A knows the size of $\underline{b}$, $\hat{b}$, and $\bar{b}$, and knows how likely it is that b equals each of these values, but does not know, in any given round of proposing the terms, which value b takes on. Therefore, A's uncertainty over B's private capabilities results in three separate values of A's probability of winning a battle p_t. Let $\underline{p}_t = \dfrac{m_A}{m_A + m_{Bt} + \bar{b}}$, $\hat{p}_t = \dfrac{m_A}{m_A + m_{Bt} + \hat{b}}$, and $\overline{p}_t = \dfrac{m_A}{m_A + m_{Bt} + \underline{b}}$, where $\underline{p}_t < \hat{p}_t < \overline{p}_t$.

A can readily infer that, in general terms, B accepts the proposal if the proposed terms are at least as good as what B expects to receive from fighting a battle: $1 - x_t \geq 1 - p_t - c_B$, or $x_t \leq p_t + c_B$. This gives a unique acceptance rule for each type. The strong B accepts if $x_t \leq \underline{x}_t$; the medium type if $x_t \leq \hat{x}_t$, and the weak type if $x_t \leq \overline{x}_t$.[12] Note that $\underline{x}_t < \hat{x}_t < \overline{x}_t$ ensures that the weak type accepts any agreement that the medium or strong type would regard satisfactory, and the medium type accepts any agreement that the strong type would regard satisfactory. Clearly, the bargaining interaction hinges on A's choice of x_t. If A proposes the reservation value of the strong B, by construction, the strong B has no reason to reject. Furthermore, the medium and weak B are both offered better terms than B would accept after fighting a battle, therefore the weak and medium B are certainly going to accept as well.

[12] Recall that A controls the prize, therefore the greater the portion of the prize A believes she can keep, the weaker A expects B to be.

What x_t will A offer? Recall that in general terms, B accepts as long as $x_t \leq p_t + c_B$. We can characterize A's optimal choice of x_t by comparing A's belief that B is weak ϕ_w to each of three possible quantities of $\hat{\phi}_w$, which A infers based on her belief about b. I will refer to these three values of $\hat{\phi}_w$ as "cutpoints," $\hat{\phi}_w^{k_1}$, $\hat{\phi}_w^{k_2}$, and $\hat{\phi}_w^{k_3}$. Cutpoint $\hat{\phi}_w^{k_1}$ determines whether A prefers setting $x_t = \underline{x}_t$ to $x_t = \hat{x}_t$. Cutpoint $\hat{\phi}_w^{k_2}$ specifies whether A prefers $x_t = \underline{x}_t$ to $x_t = \overline{x}_t$. Cutpoint $\hat{\phi}_w^{k_3}$ determines whether A prefers $x_t = \hat{x}_t$ to $x_t = \overline{x}_t$.

$\hat{\phi}_w^{k_1}$ is determined by setting the expected payoff from setting $x_t = \underline{x}_t$ greater than or equal to the expected payoff from $x_t = \hat{x}_t$. The payoff of setting $x_t = \underline{x}_t$ is $\underline{p}_t + c_B$, because it is accepted with certainty by all types as this is the most generous offer that A could make, it is not weighted by a probability. The payoff of setting $x_t = \hat{x}_t$ is $(\phi_w + \phi_m)(\hat{x}_t) + (1 - \phi_w - \phi_m)(\underline{p}_t - c_A)$. That is, A's payoff is a lottery over two possible outcomes (B accepts $\hat{x}_t$ or B rejects and A fights) weighted by their probability of occurring. Since $\hat{x}_t$ is accepted by both the weak and moderate type, the probability that A receives $\hat{x}_t$ is the sum of the probabilities that B takes on each of these two types. Since $\hat{x}_t$ is rejected only by the strong type, the value of p_t used for the battle payoff assumes B to be strong. That doesn't mean that A knows B to be strong *ex ante*, only that A's expected utility calculation reflects the understanding that if A ends up fighting B, which will happen some of the time, it must mean that in that case, B is strong. This occurs with probability $1 - \phi_w - \phi_m$. A thus solves the following inequality for ϕ_w, which, in the case of the first cutpoint, is $\hat{\phi}_w^{k_1}$:

$$\underline{p}_t + c_B \leq (\phi_w + \phi_m)(\hat{p}_t + c_B) + (1 - \phi_w - \phi_m)(\underline{p}_t - c_A)$$

which results in :

$$\hat{\phi}_w^{k_1} \leq \frac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B} - \phi_m \tag{2.1}$$

$\hat{\phi}_w^{k_2}$ is found by comparing the payoff of $\underline{x}_t$ to $\overline{x}_t$. The payoff for $\underline{x}_t$ is $\underline{p}_t + c_B$, accepted by all types. Recall that $\overline{x}_t$ is the least generous offer that A could make and it is only accepted by the weak type. The payoff for setting x_t equal to $\overline{x}_t$ is again a lottery over the possible outcomes, weighted by their probability of occurring: $\phi_w(\overline{p}_t + c_B) + \phi_m(\hat{p} - c_A) + (1 - \phi_w - \phi_m)(\underline{p} - c_A)$. That is, A's payoff is a lottery over three possible outcomes (weak B accepts $\overline{x}_t$ or medium B rejects and A fights or strong B rejects and A fights) weighted by their probability of occurring. Since $\overline{x}_t$ is accepted only by the weak type, the probability that A receives $\overline{x}_t$ is ϕ_w. Since $\overline{x}_t$ is rejected by the medium type, the value $\hat{p}_t$ is used for the battle payoff that assumes B to be medium. This occurs with probability ϕ_m. Since $\overline{x}_t$ is also rejected by the strong type, the value $\underline{p}_t$ is used for the battle payoff that assumes B to be strong. This occurs with probability $1 - \phi_w - \phi_m$. Thus, two different types of B fight if B rejects. A thus solves the following inequality for ϕ_w, which, in the case of the second cutpoint, is $\hat{\phi}_w^{k_2}$:

$$\underline{p}_t + c_B \leq \phi_w(\overline{p}_t + c_B) + \phi_m(\hat{p} - c_A) + (1 - \phi_w - \phi_m)(\underline{p} - c_A),$$

which simplifies to:

$$\hat{\phi}_w^{k_2} \leq \frac{c_A + c_B}{\overline{p}_t - \underline{p}_t + c_A + c_B} - \phi_m \frac{\hat{p}_t - \underline{p}_t}{\overline{p}_t - \underline{p}_t + c_A + c_B} \tag{2.2}$$

$\hat{\phi}_w^{k_3}$ is found by comparing the payoff of $\hat{x}_t$ to $\overline{x}_t$. As described above, the payoffs for those two strategies represent two lotteries over possible outcomes weighted by their probabilities of occurring. On the left hand side, the lottery is over whether a weak or moderately strong B accepts $\hat{x}_t$ or strong B rejects and A fights a strong B. On the right hand side, the lottery is over whether a weak B accepts $\overline{x}_t$ or a moderately strong B rejects and A fights or a strong B rejects and A fights. A solves the inequality below with respect to ϕ_w, which in this case represents $\hat{\phi}_w^{k_3}$:

$$(\phi_w+\phi_m)(\hat{p}_t+c_B)+(1-\phi_w-\phi_m)(\underline{p}_t-c_A) \leq \phi_w(\overline{p}_t+c_B)+\phi_m(\hat{p}-c_A)+(1-\phi_w-\phi_m)(\underline{p}-c_A),$$

which simplifies to:

$$\hat{\phi}_w^{k_3} \leq \phi_m \frac{c_A+c_B}{\overline{p}_t-\underline{p}_t} \tag{2.3}$$

In any given circumstance, there are only two cutpoints that matter. But which two matter varies. If A prefers the most generous offer $\underline{x}_t$ to the moderately generous offer $\hat{x}_t$, then it is irrelevant whether A prefers the least generous offer $\overline{x}_t$ to $\hat{x}_t$. However, knowing that A prefers $\underline{x}_t$ to $\hat{x}_t$ does not tell us whether A prefers the most generous offer $\underline{x}_t$ to the least generous offer $\overline{x}_t$, because the payoffs from accepted/rejected offers are weighted by the probabilities of their occurring, and, therefore, they are lotteries. Similarly, if A prefers $\hat{x}_t$ to $\underline{x}_t$, then it is irrelevant whether A prefers $\overline{x}_t$ to $\underline{x}_t$. In contrast, whether A prefers $\hat{x}_t$ to $\overline{x}_t$ is critical for A to make a decision about the size of x_t.

Once A knows the value of each cutpoint, A compares her exogenous belief ϕ_w to these quantities. Recall that cutpoint $\hat{\phi}_w^{k_1}$ determines whether A prefers the most generous offer $\underline{x}_t$ to the moderately strong offer $\hat{x}_t$, while cutpoint $\hat{\phi}_w^{k_2}$ specifies if A prefers $\underline{x}_t$ to the least generous offer $\overline{x}_t$. When A's belief about B being weak is less than or equal to the lower of these cutpoints, $\phi_w \leq min\{\hat{\phi}_w^{k_1}, \hat{\phi}_w^{k_2}\}$, A proposes the terms of $x_t = \underline{x}_t$, which all three types accept, thus ensuring peace. Cutpoint $\hat{\phi}_w^{k_3}$ determines whether A prefers moderately strong offer $\hat{x}_t$ to the least generous offer $\overline{x}_t$. When A's belief about B being weak is greater than the higher value of cutpoints $\hat{\phi}_w^{k_2}$ and $\hat{\phi}_w^{k_3}$, $\phi_w > max\{\hat{\phi}_w^{k_2}, \hat{\phi}_w^{k_3}\}$, A proposes the least generous terms $x_t = \overline{x}_t$, which only the weakest type accepts. Finally, when A's belief about B being weak is greater than the first cutpoint and is less than the third cutpoint, $\hat{\phi}_w^{k_1} < \phi_w < \hat{\phi}_w^{k_3}$, A sets $x_t = \hat{x}_t$, which both the weak and medium types accept.

If A sets $x_t = \underline{x}_t$, all three types accept.

If A sets $x_t = \hat{x}_t$ in period t, then in $t+1$, A either knows that B cannot be strong (if B accepted in t), or that B must be strong (if B rejected). If B accepted then A knows that the only two possible types left are the weak and moderate. A can either choose the least generous offer $\overline{x}_t$, which risks war in the case that B is moderately strong, or $\hat{x}_t$, which is accepted by both types. To make the choice of x_t, A compares her payoff from receiving $\hat{x}_t$ with certainty on the one hand to receiving $\overline{x}_t$ weighted by the probability that B is weak ϕ_w and receiving a war payoff with probability $1 - \phi_w$, or $\hat{x}_t \leq \phi_w(\overline{x}_t) + (1 - \phi_w)(\hat{p}_t - c_A)$. Once A solves this inequality for ϕ_w, which is now A's updated belief that B is weak, A keeps x_{t+1} equal to $\hat{x}_{t+1}$, if:

$$\phi_w' \leq \frac{c_A + c_B}{\overline{p}_t - \hat{p}_t + c_A + c_B} \tag{2.4}$$

Otherwise, A proposes $\overline{x}_{t+1}$ and risks war if B is moderate.

If A sets x_t to $\overline{x}_t$ in period t, then in $t+1$, A either knows that B must be weak (if B accepted in t), or that B is either moderately strong or strong (if B rejected in t). If B rejected, A can further update her beliefs based on the outcome of the previous battle. Battles are won probabilistically based on the value of p_t. If A won, A believes that B is moderately strong with probability $\phi_m^{won} = \dfrac{\hat{p}_t\phi_m}{\hat{p}_t\phi_m + \underline{p}_t(1-\phi_m)}$, which describes A's probability of victory when B is moderately strong as a proportion of A's total probability of winning, i.e. the sum of A's probability of winning given B is moderately strong and A's probability of winning given B is strong, all weighted by their probabilities of occurring. Since B can now only be either strong or moderately strong, $\phi_m \leq 1$. If A lost, A believes that B is moderately strong with probability $\phi_m^{lost} = \dfrac{(1-\hat{p}_t)\phi_m}{(1-\hat{p}_t)\phi_m + (1-\underline{p}_t)(1-\phi_m)}$, which describes B's probability of winning if B is moderately strong as a proportion of the sum of B's probabilities of winning when B is moderately strong and strong, all weighted by their probabilities of occurrence. To make the choice of x_t, A compares her payoff from receiving $\underline{x}_t$ with certainty to receiving $\hat{x}_t$ weighted by the probability that B is moderately strong ϕ_m and receiving a war payoff with probability $1-\phi_m$, or $\underline{x}_t \leq \phi_m(\hat{x}_t) + (1-\phi_m)(\underline{p}_t - c_A)$, where ϕ_m equals either ϕ_m^l or ϕ_m^w and the war payoff is based on A fighting a strong B. Once A solves this inequality for ϕ_m', which is now A's updated belief about B being moderately strong, A sets x_{t+1} equal to $\underline{x}_{t+1}$, if:

$$\phi'_m \leq \frac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B} \tag{2.5}$$

Otherwise, A proposes the moderate terms $\hat{x}_{t+1}$.

Major takeaways about the three types setup: The inequalities describing A's decision-making process with respect to determining the initial value of x in period t (2.1, 2.2, 2.3) and the inequalities describing A's follow-up decisions to determine the value of x in period $t+1$ (2.4, 2.5) have common features. First, for each of these five inequalities the value of the right-hand side increases as the sum of the costs that the sides pay when they fight grows. Second, the right-hand side increases when the difference between A's probability of defeating B when B is relatively weaker and when B is relatively stronger shrinks. In other words, the larger the difference the more likely A is to risk war, as the difference increases, the denominator gets larger, and the fraction gets smaller, and more and more values of cutpoints (2.1-2.3) or updated beliefs (2.4, 2.5) will lie above the threshold. *The probability of A proposing the terms that risk war grows as fighting becomes less costly and as the difference between A's probability of defeating two different types of B increases.*

In addition, note that since the difference between A's probability of winning when B is relatively weaker compared to winning when B is relatively stronger is the smallest when m_A and m_B are relatively equal (though note that it's not exactly when they are equal because of b), this means that parity is associated with a decrease in ϕ, which means a larger range of values of ϕ encourage A to risk war. This relationship between parity and war is consistent with the empirical and theoretical literature on the distribution of power and war onset (Reed 2003).

Evaluation of the status quo: Recall that π denotes the distribution of benefits in the dyad and reflects how much of the prize side A controls. Both sides of a dyad may want to challenge status quo if they do not benefit from the distribution of the good proportionally to their relative power. B is dissatisfied with status quo and wants to challenge A if B's share of benefits is less than B's payoff from fighting, $1 - \pi < 1 - p_t - c_B$ or $\pi > p_t + c_B$, using the value of p_t that B knows from nature. Similarly, A is dissatisfied and wants to challenge status quo if she believes that she could receive more of the prize from fighting: $\pi < p_t - c_A$. Recall that A has three estimates of her probability of winning $\underline{p}_t$, $\hat{p}_t$ and $\overline{p}_t$. Thus, in the simulation, A first determines the value of the potentially proposed terms x_t, and then compares π to the estimate of p_t that implies A's best guess about her probability of winning militarily. If $x_t = \underline{x}_t$, then A compares π to $\underline{p}_t$; if $x_t = \hat{x}_t$, then A compares π to $\hat{p}_t$; if $x_t = \overline{x}_t$, then A compares π to $\overline{p}_t$.

Sequence of events: Once either side is dissatisfied, A starts making offers. B accepts A's first offer if A proposes $\underline{x}_t$, which is accepted by all three types. The new settlement is recorded as $\pi = p_t + c_B$. This reflects the sides resolving their conflict of interests without resorting to arms. If A proposes $\hat{x}_t$, B accepts if he is weak or moderately strong. If after B's acceptance, A believes that B is moderately strong, A keeps $\hat{x}_t$ and the new π is recorded. If after B's acceptance, A believes that B is weak, A proposes $\overline{x}_t$ and risks a war if B is moderately strong. Finally, if A proposes $\overline{x}_t$ and B accepts, B is weak. If B rejects, they fight a battle and A updates her belief about B being moderately strong based on the battle outcome. If A believes that B is strong, after the first battle, A offers $\underline{x}_{t+1}$, which settles the conflict. In contrast, offering $\hat{x}_{t+1}$ based on the updated belief that B is moderately strong risks another battle if B is strong. Either in period t or in the following period $t + 1$, A makes a generous enough offer that B accepts and the sides *settle*. In other words, the new settlement π is recorded each time when the sides settle and uncertainty over B's type is solved.

Peace phase: Recall Powell's (1996, 1999, 2004) argument that the probability of war onset is proportional to the gap between the distribution of power and the distribution of benefits. I suggest that Powell's idea is applicable to defining the probability of war recurrence as well. Ceasefire agreements are documents that can be seen as codifying the distribution of benefits between the former belligerents. The more closely the distribution of benefits in the document reflects the sides' expectations about the next potential battle (i.e. the distribution of power), the longer the ceasefire will survive (assuming away random rapid shifts in power).

The null model's setup allows the belligerents to settle only after A learns (through bargaining and fighting) the true type of B and B accepts only that offer, which reveals B's private capabilities. So, every ceasefire of the benchmark model establishes a political settlement that reflects a "true" distribution of power in the dyad. The greater the gap between the distribution of power and the distribution of benefits institutionalized by the ceasefire settlement, the more severe the conflict of interests remains. One can think of the severity of the remaining conflict of interests as the baseline fragility of a ceasefire agreement. Since the null model excludes any third-party behavior, the sides are never pressured to settle prematurely, therefore, the gap between the distribution of power and the distribution of benefits will be absent to minimal in the null model at the moment of signing the settlement.

In the null model, it is only due to exogenous shocks in capabilities that war could potentially recur. The rate of conflict recurrence in the null model is expected to be much lower than the one observed empirically.

In the peace phase, A does not make explicit offers to B until either side is dissatisfied. One can think of this period of peace duration as if A's last proposal is implied and remains the same from round to round. As long as B accepts that implicit offer, the ceasefire lasts. Once, due to exogenous shock, either side's capabilities grow and that side could potentially gain from revising the settlement (condition that $\pi < p_t + c_B$ for A or $\pi > p_t + c_B$ for B is true), A makes explicit offers. If A's first offer is accepted by B then the settlement value π is updated to reflect the new proportion of the good that A holds. If A's first offer is rejected, the peace phase ends and the sides fight a battle. In the null model, the war phase lasts the maximum of two battles as I have assumed only three possible types of B.

As discussed above, the bargaining arguments about war as a process explain that war may allow for learning through revealing capabilities by fighting battles. Therefore, peace settlements signed at the conclusion of information-driven wars should reflect the overlapping expectations about what exact division of benefits the sides are willing to uphold. To borrow an argument from the international economy paper by Downs, Rocke and Barsoom (1996), in the world with no external pressure to settle prematurely, we should observe high rates of compliance with peace settlements, as the belligerents do not have to sign agreements that are not self-enforcing. Downs, Rocke and Barsoom (1996) argue that compliance with institutions does not speak to the strength of institutions but rather the fact that those institutions are endogenous to the preceding bargaining process, which in my model is represented by the process of fighting battles.[13]

Summary: In a world without third parties, conflicts of interests may turn into wars if there is uncertainty over capabilities. Specifically, a war starts if the offer-making side cannot estimate correctly the power of the offer-rejecting/accepting side. Observing opponent's decisions to reject/accept and fighting battles allows the offer-making side to revise her expectations about the respective probabilities of winning the next potential battle. Eventually, such updating leads to the overlap in the sides' expectations and a ceasefire settlement emerges. If the sides settle only once their expectations converge, the peace settlement is reflective of the true distribution of power in the dyad. Therefore, such "organic" peace agreements are void of conflicts of interests, unless exogenous shock to relative capabilities leads to the challenger's dissatisfaction with the status quo. On average, peace settlements are expected to be stable in the null model, and more stable than what is observed empirically. This is an important takeaway of the null model: it does not represent a conflict recurrence rate realistically. By introducing third-party behavior in the extended model (chapter 4), I demonstrate that conflict recurrence becomes more common due to premature settlements and extended support after the conflict is over. Therefore, an extended model simulates the international system in a more realistic fashion and underscores the importance of taking third-party behavior into account.

[13] Fearon (1998) and Von Stein (2005) have made similar arguments about states "selecting" themselves into those institutions that they want to uphold.

2.2 Pseudocode for the Null Model

2.2.1 Big picture outline

The following sequence summarizes the repeating loop (steps 2-5) of the simulation:

1. Initialize the world, set up the agents and the dyads. For each dyad, B is assigned a type (weak, moderately strong, or strong) randomly.
2. For each dyad, check whether either side is dissatisfied with the status quo. If both are satisfied, the dyad stays at peace. If either side is dissatisfied, the dyad's status is updated to having a conflict of interests.
3. For each dyad that has a conflict of interests: A determines which terms to propose based on her beliefs and the difference between defeating two different types of B.
    - If A proposes the terms that all three types accept, the dyad's status is updated to peaceful.
    - If A proposes the terms that only moderately strong and weak types accept:
        (a) If B is strong, B rejects. A's beliefs are reset. Uncertainty solved. The dyad's status is updated to armed conflict,
        (b) If B is moderately strong or weak, B accepts. Uncertainty remains. Determine if A's belief that B is weak encourages A to risk war:**??**
            i. If no, the terms are set to reflect that B is moderately strong; the dyad's status is updated to peaceful.
            ii. If yes, A proposes the terms that only weak B would accept:
                A. If B accepts, B is weak, the beliefs and terms are reset to reflect this; the dyad's status is updated to peaceful.
                B. If B rejects, B is moderately strong, the beliefs are reset to reflect this, uncertainty is solved; the dyad's status is updated to armed conflict.
    - If A proposes the terms that only a weak type would accept:
        (a) If B is weak, B accepts and the dyad's status is set to peaceful.
        (b) If B is moderately strong or strong, B rejects, uncertainty is not solved; the dyad's status is reset as armed conflict.
4. For each dyad that is at armed conflict:
    - If uncertainty is solved. The sides fight one battle.
        (a) If B wins, the terms are set to reflect that B is of a stronger type.
        (b) If A wins, the terms are set to the last offer, however given A's updated beliefs, the dyad will not escalate to armed conflict next period.

- If uncertainty is not solved. Uncertainty is over whether B is moderately strong or strong. The sides fight a battle. If A wins, her belief that B is moderately strong increases. If B wins, A's belief that B is moderately strong decreases. A compares her updates belief to the utility of proposing moderately strong v. strong terms.
    (a) If A estimates that B is moderately strong, A proposes the terms that only a moderately strong B accepts.
        i. If B is moderately strong, he accepts, the dyad status is peaceful.
        ii. If B is strong, he rejects; A's beliefs are updated to reflect that. Uncertainty is solved. The dyad's status stays at armed conflict.
    (b) If A estimates that B is strong, A proposes the terms that any type of B accepts. The terms are updated.

5. Go to step 2.

2.2.2 Detailed pseudocode

The following sequence of steps is written in the Java programming language with the use of RePast libraries to reflect the discussion above.

I. **Do these steps once:**

1. Create a `numstates` number of political actors in the "world" (minimum of 2), each characterized by randomly generated $m_i \in [0, 1]$.
2. All actors are paired up in non-directional dyads and are assigned randomly either the role of the potential defender A, or the role of the potential challenger B.
3. All dyads are located on the `ArrayOfPeacefulDyads`.
4. Each dyad is assigned a randomly generated value π that reflects the division of the disputed good in the dyad. π is the proportion of the good controlled by A.
5. Each dyad is assigned two randomly generated beliefs for A: ϕ_w and ϕ_m, such that $\phi_w + \phi_m \leq 1$.
6. Each dyad is assigned three randomly generated values of b for B, such that $\underline{b} < \hat{b} < \overline{b}$ AND $b \leq (0.2 \times m_B)$.
7. Set `termsDeterminedInPreviousRound`=FALSE for each dyad.

To ensure that events take place sequentially, execute for-loop 1 and for-loop 2 if the tick is even, execute for-loop 3 if the tick is odd.

II. **For-loop 1:** for each dyad located on the `ArrayOfPeacefulDyads`, each tick follows this sequence:

1. Determine whether a conflict of interests exists. In each dyad, B checks whether $\pi > p_t + c_B$ (B knows which type he is, so the choice of p_t is obvious for B) and A checks whether $\pi < p_t - c_A$. With respect to the choice of p_t, A first calculates, $\phi_w^{k_1} = \frac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B} - \phi_m$, $\hat{\phi}_w^{k_2} = \frac{c_A + c_B}{\overline{p}_t - \underline{p}_t + c_A + c_B} - \phi_m \frac{\hat{p}_t - \underline{p}_t}{\overline{p}_t - \underline{p}_t + c_A + c_B}$, $\hat{\phi}_w^{k_3} = \phi_m \frac{c_A + c_B}{\overline{p}_t - \underline{p}_t}$. When $\phi_w \leq min\{\hat{\phi}_w^{k_1}, \hat{\phi}_w^{k_2}\}$, A uses $\underline{p}_t$; when $\phi_w > max\{\hat{\phi}_w^{k_2}, \hat{\phi}_w^{k_3}\}$, A uses $\overline{p}_t$; when $\hat{\phi}_w^{k_1} < \phi_w < \hat{\phi}_w^{k_3}$, A uses $\hat{p}_t$.
   i. If either condition holds, the dyad is added to the `ArrayOfConflictsOfInterest`.

III. **For-loop 2:** for each dyad located on the `ArrayOfConflictsOfInterest`, each tick follows this sequence:

1. If `termsDeterminedInPreviousRound`=FALSE, determine x_t. When $\phi_w \leq min\{\hat{\phi}_w^{k_1}, \hat{\phi}_w^{k_2}\}$, A sets $x_t = \underline{x}_t$, when $\phi_w > max\{\hat{\phi}_w^{k_2}, \hat{\phi}_w^{k_3}\}$, A sets $x_t = \overline{x}_t$, when $\hat{\phi}_w^{k_1} < \phi_w < \hat{\phi}_w^{k_3}$, A sets $x_t = \hat{x}_t$.
   i. If A proposes terms $\underline{x}_t$:
      A. Every type of B accepts.
         • Update $\pi = \underline{p}_t + c_B$. Move the dyad to the `ArrayOfPeacefulDyads`.
   ii. If A proposes terms $\hat{x}_t$:
      A. If B is strong, he rejects.
         • Move the dyad to the `ArrayOfArmedConflicts`. Set boolean `uncertaintySolved` as true. Since only a strong type rejects, update A's beliefs as: $\phi_w = 0.05, \phi_m = 0.05$.
      B. If B is weak or moderately strong, he accepts.
         • if $\phi_w \leq \frac{c_A + c_B}{\overline{p}_t - \hat{p}_t + c_A + c_B}$:

         A keeps $\hat{x}_t$. Move the dyad to the `ArrayOfPeacefulDyads`. A's beliefs are updated to reflect that A knows that B is not strong, so the belief that B is strong is reduced to $(1 - \phi_w - \phi_m) = 0.05$, the rest of the difference is distributed between ϕ_m and ϕ_w equally.
         • if $\phi_w > \frac{c_A + c_B}{\overline{p}_t - \hat{p}_t + c_A + c_B}$:

         The conflict of interest persists, as A now thinks it is possible to keep $\overline{x}$ in the following round. Keep the dyad on this array and set `termsDeterminedInPreviousRound` as TRUE.
   iii. If A proposes terms $\overline{x}_t$ and `termsDeterminedInPreviousRound` is FALSE:
      A. If B is weak, he accepts.
         • Update $\pi = \overline{p}_t + c_B$. Update A's beliefs $\phi_w = 0.99$ and $\phi_m = 0.05$ to reflect that A knows that B is weak. Move the dyad to the `ArrayOfPeacefulDyads`.
      B. If B is moderately strong or strong, he rejects.
         • Move the dyad to the `ArrayOfArmedConflicts`. Set boolean `uncertaintySolved` as FALSE.

2. If `termsDeterminedInPreviousRound`=TRUE, do:
   i. A proposes terms $\overline{x}_t$ and `termsDeterminedInPreviousRound` is TRUE
      A. If B is weak, he accepts.
         • Update $\pi = \overline{p}_t + c_B$. Update A's beliefs $\phi_w = 0.99$ and $\phi_m = 0.05$ to reflect that A knows that B is weak. Move the dyad to the `ArrayOfPeacefulDyads`.
      B. If B is moderately strong, he rejects.
         • Move the dyad to the `ArrayOfArmedConflicts`. Set boolean `uncertaintySolved` as TRUE. Since only moderately strong B would reject that offer, A knows B's type now, update beliefs such that $\phi_w = 0.05, \phi_m = 0.99$

IV. **For-loop 3:** for each dyad located on the `ArrayOfArmedConflicts`, each tick does:

1. The sides fight a battle. Generate a random value p' that is drawn from a uniform distribution, $p' \in [0, 1]$.
2. If `uncertaintySolved` is TRUE and B rejected $\hat{x}_{t-1}$:
   i. If $p' \leq \hat{p}_t$, A wins the battle, A keeps $\hat{x}_t$; update $\pi = \hat{p}_t + c_B$; move the dyad to the `ArrayOfPeacefulDyads`.

   ii. If $p' > \hat{p}_t$, A loses, she keeps $\underline{x}_t$; update $\pi = \underline{p}_t + c_B$; move the dyad to the `ArrayOfPeacefulDyads`.
3. If `uncertaintySolved` is TRUE and B accepted $\hat{x}_{t-1}$, but rejected $\overline{x}_{t-1}$:
   i. If $p' \leq \hat{p}_t$, A wins the battle, A keeps $\overline{x}_t$; update $\pi = \overline{p}_t + c_B$; move the dyad to the `ArrayOfPeacefulDyads`.

   ii. If $p' > \hat{p}_t$, A loses, she keeps $\hat{x}_t$; update $\pi = \hat{p}_t + c_B$; move the dyad to the `ArrayOfPeacefulDyads`.
4. If `uncertaintySolved` is FALSE (B rejected $\overline{x}_{t-1}$):
   i. If $p' \leq \hat{p}_t$, A wins a battle. A calculates updated belief about the value of ϕ_m: calculates $\phi_m^w = \dfrac{\hat{p}_t \phi_m}{\hat{p}_t \phi_m + \underline{p}_t (1 - \phi_m)}$. A then compares the new belief about ϕ_m to the value that is derived from comparing A's utility of proposing the terms as if B is moderately strong, as opposed to proposing the terms as if B is strong:

      A. If $\phi_m^w \leq \dfrac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B}$. A sets x_{t+1} equal to $\underline{x}_{t+1}$.
         • B accepts. Update $\pi = \underline{p}_t + c_B$. A's beliefs are updated as $\phi_w = 0.05$ (to reflect that A has found out that B is not weak), and $\phi_m = \phi_m^w$ (to

reflect the updated belief ϕ_m after fighting a battle).[14] Move the pair to the `ArrayOfPeacefulDyads`.

B. If $\phi^w_m > \dfrac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B}$. A sets x_{t+1} equal to $\hat{x}_{t+1}$.

- If B is moderately strong, he accepts, update $\pi = \hat{p}_t + c_b$; A's beliefs are updated to reflect that only a moderately strong type would accept the offer: $\phi_w = 0.05, \phi_m = 0.99$; move the dyad to `ArrayOfPeacefulDyads`.

- If B is strong, he rejects. A's beliefs are updated to reflect that only a strong type would reject the offer: $\phi_w = 0.05, \phi_m = 0.05$. Set boolean `uncertaintySolved` as TRUE, keep on this array.

ii. If $p' > \hat{p}_t$, B won. A calculates updated belief about the value of ϕ_m: $\phi^l_m = \dfrac{(1-\hat{p}_t)\phi_m}{(1-\hat{p}_t)\phi_m + (1-\underline{p}_t)(1-\phi_m)}$. A then compares the new belief about ϕ_m to the value that is derived from comparing A's utility of proposing the terms as if B is moderately strong, as opposed to proposing the terms as if B is strong:

A. If $\phi^l_m \leq \dfrac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B}$. A sets x_{t+1} equal to $\underline{x}_{t+1}$.

- B accepts. Record $\pi = \underline{p}_t + c_B$. A's beliefs are updated as $\phi_w = 0.05$ (to reflect that A has found out that B is not weak), and $\phi_m = \phi^w_m$ (to reflect the updated belief ϕ_m after fighting a battle).[15] Move the dyad to the `ArrayOfPeacefulDyads`.

B. If $\phi^l_m > \dfrac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B}$. A sets x_{t+1} equal to $\hat{x}_{t+1}$.

- If B is moderately strong, he accepts, update $\pi = \hat{p}_t + c_B$. A's beliefs are updated to reflect that only a moderately strong type would accept this offer: $\phi_w = 0.05, \phi_m = 0.99$. Move the dyad to `ArrayOfPeacefulDyads`.

- If B is strong, he rejects. A's beliefs are updated to reflect that only a strong type would reject the offer: $\phi_w = 0.05, \phi_m = 0.05$. Set boolean `uncertaintySolved` as TRUE. Keep the dyad on this array.

[14] Even though A has won, the most generous offer of $\underline{x}_{t+1}$ reflects that A's belief that B is strong is the highest, therefore, the dyad is extremely unlikely to return to the conflict of interests stage, and if it does, A clearly believes that B is most likely to be strong and therefore will again make an appeasing offer.

[15] B has won, the most generous offer of $\underline{x}_{t+1}$ reflects that A's belief that B is strong is the highest, therefore, the dyad is extremely unlikely to return to the conflict of interests stage, and if it does, A clearly believes that B is most likely to be strong and therefore will again make an appeasing offer.

2.3 Analysis of the Computer Simulations' Numerical Output

2.3.1 Illustrative runs of the benchmark model

There are 200 agents in the artificial international system, paired up randomly into 50 dyads (these are default values, the range of possible values is shown in Table 2.1). The agents are paired up randomly and are assigned either a role of side A that has the power to propose the terms of the deal or the role of side B that may reject or accept the terms proposed by A. Over the course of the simulation, a dyad may either be at peace, or experience a conflict of interests, or fight a violent conflict. The four panels in Figure 2.1 show the time series of counts of dyads at peace (shown in green), at potential conflict (depicted in blue), and at armed conflict (shown in red) for the simulation runs with 15, 50, 75, and 100 dyads. As these graphs demonstrate, the number of dyads in the artificial international system does not impact the major trend: initially, some proportion of dyads experiences conflicts of interests, and some of those disputes become violent. As time goes on (after the initial 10-17 time units), the conflicts of interests no longer arise, as they get settled through redistributing the benefits in accordance with the distribution of relative power in each dyad through A's learning of B's type.

Figures 2.2 and 2.3 show that the process of "learning" (i.e., updating A's beliefs) takes place in the initial ticks of the simulation and the subsequent interactions become peaceful as a result of this updating. In particular, the graphs in Figure 2.2 chart the differences between A's beliefs about B's strength in time period t from A's beliefs in time period $t-1$ (i.e., $\phi_t - \phi_{t-1}$). After the initial 10-17 time units, the "learning" of B's type is complete. It is important to note that A's belief that B is strong rarely decreases, as wars break out due to A's underestimation of B's strength (as demonstrated by the bottom panel in Figure 2.2). The only scenario that allows for decrease in A's belief that B is not as strong as previously thought is described in step **??** of the pseudocode. Figure 2.3 includes three graphs that show the absolute differences between A's beliefs about B's strength in time t relative to time $t-1$. Similarly, the graphs in Figure 2.4 chart the changes in settlement π over the simulation time. While the top panel in Figure 2.4 plots the values of each dyad's division of the disputed good over the simulation time, the bottom panel shows the absolute differences in settlement time t relative to time $t-1$, i.e., $\pi_t - \pi_{t-1}$.

Figures 2.1-2.4 describe the same key feature of the null model that sets it apart from the extended model: the established peace treaties are stable, when (a) the sides settle when B's type is revealed and (b) there are no shifts in capabilities after a peace deal is reached. In the extended model, I demonstrate that some conflicts are ended prematurely (i.e. without revealing B's type) due to power mediators' pressure. Furthermore, after a ceasefire, support-providing third parties sometimes generate shifts in capabilities that may either recreate or remove uncertainty about relative probability of victory. It is important to note that shifts in relative capabilities may occur due to factors other than external support: discovery of natural resources or acquiring of nuclear weapons may increase military capabilities of a government without receiving outside support. While I do not model "natural" shifts in capabilities in this project, this should be an important future extension.

2.3.2 Scenarios that highlight the importance of initial beliefs

In each dyad the sides decide whether they would like to challenge the status quo. If either side is dissatisfied with the current distribution of benefits, then a conflict of interests arises, as one of the sides believes that s/he could receive more benefits by fighting a military contest. I now describe three groups of scenarios that may unfold in the null model, under the condition that one of the sides is dissatisfied with the distribution of benefits in the dyad. All cases in the model fit into one of the following seven scenarios:

1. A sets x_t to $\underline{x}_t$, and all types of B accept.
2. A proposes $\hat{x}_t$ and
   (a) a moderately strong or a weak type B accepts, A decides not to risk war,
   (b) a moderately strong or a weak type B accepts, A decides to risk war (thinking that B may be weak),
   (c) a strong B rejects and the sides fight one battle.
3. A sets x_t to $\overline{x}_t$ and
   (a) a weak B accepts,
   (b) a moderately strong B rejects in t but accepts $\hat{x}$ in $t+1$ (one battle),
   (c) a strong B rejects in t, then A proposes the terms of $\hat{x}$ in $t+1$ rejected by strong B, finally strong B accepts $\underline{x}$ in $t+2$ (two battles).

Scenario 1. Agent number 23 and agent number 81 are paired up such that 23 is assigned the role of side A and 81 is side B. I refer to these agents as A_{23} and B_{81}. Following Joyce (2008), agent capabilities in the model follow beta distribution with parameters $\alpha = 0.05$ and $\beta = 1.25$, such that the distribution of capabilities in the world approximates the empirical distribution of the Composite Index of National Capability (CINC) score (Singer, Bremer and Stuckey 1972). The comparison of the simulated capability data to the emprirical distribution of CINC scores is shown in Figure 2.5 (a). Panel (b) of Figure 2.5 presents the distribution of agent capability in a single run. The capability score $m_i \in [0, 1]$.

A_{23}'s capabilities are 0.00459 and B_{81}'s observable capabilities equal 0.02135, so even without taking B_{81}'s private capabilities into account, it is clear that B_{81} is a more powerful government in this pair. The private component b of B's capability is generated randomly such that $\underline{b} < \hat{b} < \bar{b}$. In A_{23}'s perception, there is a 0.13 chance (exogenously given ϕ_w) that B_{81} is weak. If so, B_{81}'s private capability component equals $\underline{b} = 0.00543$, so that A_{23} controls $\dfrac{m_{A_{23}}}{m_{A_{23}} + m_{B_1} + \underline{b}} = \dfrac{0.00459}{0.00459 + 0.02135 + 0.00543} = 0.14638 \approx 14.6\%$ of the dyad's total capabilities. There is a 0.48 chance that B_{81} is moderately strong (exogenously given ϕ_m), $\hat{b} = 0.00741$, which would let A_{23} control $\approx 13.8\%$ of the dyad's power. Finally, in A_{23}'s perception, there is a $(1 - \phi_w - \phi_m) = 0.39$ chance that B_{81} is strong and $\bar{b} = 0.00879$, which would allow A_{23} to control $\approx 13.2\%$ of the dyad's capabilities.

The simulation starts with all dyads being classified as peaceful. During the first time period, each dyad's participants assess whether they are satisfied with the distribution of the good recorded as π, the proportion of the good controlled by side A. Recall that side B is dissatisfied if his exogenously given proportion of the good $1 - \pi$ is less than what B could receive from fighting a war, which is $(1 - p_t) - c_B$, or $\pi > p_t + c_B$. A is dissatisfied with the status quo if she receives fewer benefits π than what A could receive from fighting $\pi < p_t - c_A$. For this dyad, A_{23}'s probability of winning (which equals the proportion of capabilities controlled by A in the dyad) against a weak B_1 is $\overline{p}_t = 0.146$; A_{23}'s chance of winning if B_{81} is moderate strong is $\hat{p}_t = 0.138$; A_{23}'s probability of winning if B_{81} is strong equals $\underline{p}_t = 0.132$. Costs of fighting a battle, c_A and c_B, were generated as 0.001 for A_{23}

and 0.015 for B_{81}. Each value of cost of fighting $\in [0, 0.1]$ and is drawn from a uniform distribution. In this dyad, π was generated equal to 0.8557, which is less than B_{81}'s payoff from fighting a battle regardless of his type. Since A_{23} enjoys many more benefits than what is proportional to her relative power, B_{81} is dissatisfied.

Even though in this example A is satisfied, it is important to note that A's decision-making process is more involved as A calculates three values that describe A's "cutpoint" values for preferring $\underline{x}_t$ to $\hat{x}_t$, $\underline{x}_t$ to $\overline{x}_t$, and $\hat{x}_t$ to $\overline{x}_t$. If A's belief that B is weak ϕ_w is less than or equal to the minimum of cutpoints 1 and 2, then A's best guess is that B is a strong type. In this example, A_{23}'s belief that B_1 is weak is $\phi_w = 0.13$. A_{23} first calculates cutpoints 1 and 2 and then compares ϕ_w to the minimum of the two values. $\phi_w^{k_1} = \dfrac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B} - \phi_m = \dfrac{0.01 + 0.015}{0.138 - 0.132 + 0.01 + 0.015} - 0.48 = 0.3264516129$, $\hat{\phi}_w^{k_2} = \dfrac{c_A + c_B}{\overline{p}_t - \underline{p}_t + c_A + c_B} - \phi_m \dfrac{\hat{p}_t - \underline{p}_t}{\overline{p}_t - \underline{p}_t + c_A + c_B} =$ $\dfrac{0.01 + 0.015}{0.146 - 0.132 + 0.01 + 0.015} - 0.48 \times \dfrac{0.138 - 0.132}{0.146 - 0.132 + 0.01 + 0.015} = 0.03226 - 0.48 \times 0.42254 =$ 0.56718. Since 0.13 is less than either of these values, A_{23} indeed believes that B_{81} is strong.

For the purpose of illustration, I calculate the values for cutpoint 3 to show that A's beliefs are exclusive of each other. If A's belief ϕ_w is greater than cutpoint 1 and is less than cutpoint 3, then A believes that B is moderately strong. In this example, $\hat{\phi}_w^{k_3} = \phi_m \dfrac{c_A + c_B}{\overline{p}_t - \underline{p}_t} = 0.48 \times \dfrac{0.01 + 0.015}{0.146 - 0.132} = 0.85714$. Since 0.13 is less than both cutpoint 1 and cutpoint 3, A_{23} does not believe that B_1 is moderately strong. Finally, if A's belief ϕ_w is greater than the maximum of cutpoints 2 and 3, then A believes that B is weak. Again, A_{23}'s belief is strictly less than any of the cutpoints, therefore A_{23} believes that B_{81} is strong. Comparing π to the proportion of the good that A_{23} would need to give up to in the case of fighting a war against B_{81}, $\underline{p}_t + c_B$, A_{23} determines $0.8557 > (0.132 + 0.015)$. This means that A_{23} receives more benefits than what is proportional to her relative power in the dyad and is therefore satisfied with the status quo (unlike B_{81}). See section II.1 of the pseudocode section for a general description of peaceful interactions.

In time period 2, the dyad is classified as having a conflict of interest because of B_{81}'s dissatisfaction with the status quo. A_{23} proposes the new terms of political agreement to B_{81}. Side A needs to determine whether to set the terms x_t to $\underline{x}_t, \hat{x}_t$, or $\overline{x}_t$ based on her beliefs about B's type. As determined above, A_{23} believes that B_{81} is strong, therefore A_{23} proposes to keep $\underline{x}_t = \underline{p}_t + c_B = 0.132 + 0.015 = 0.147$, as opposed to the old division of benefits $\pi = 0.8557$. Since $\underline{x}_t$ is the lowest proportion of benefits that A_{23} could hope to control in this dyad, any type of B accepts the new terms as this is exactly what the strong B would consider as a fair distribution of benefits, yet, the moderately strong B and weak B are receiving more benefits than they would otherwise, so they certainly accept as well. The status quo is updated $\pi = 0.147$. A_{23}'s beliefs about B_{81}'s type do not change due to this interaction. The dyad's status is reset as peaceful.

In time period 3 and all subsequent rounds of the simulation, A_{23} and B_{81} determine whether they are dissatisfied with the status quo of $\pi = 0.147$. Given that the sides' capabilities do not change and the cost of fighting stays the same, both A_{23} and B_{81} will stay satisfied with the status quo, so the status of the dyad will remain peaceful forever in this artificial international system.

Scenario 2a. Let beliefs be $\phi_w = 0.15$ and $\phi_m = 0.7$ for the described above pair of A_{23} and B_{81}. These beliefs generate $\hat{\phi}_w^{k_1} = 0.106$, $\hat{\phi}_w^{k_2} = 0.533$, and $\hat{\phi}_w^{k_3} = 1.25$, which means that $\hat{\phi}_w^{k_1} < \phi_w < \hat{\phi}_w^{k_3}$. Given these beliefs, A_{23} proposes terms $\hat{x}_t = \hat{p}_t + c_B = 0.138 + 0.015 =$

0.153 in period 1. This offer is greater than what a weak type of B would accept (scenario 1), so a weak B accepts. Similarly, a moderately strong type accepts, as this offer reflects what a moderately strong type would receive after fighting a battle. If B_{81} happens to be strong, then B_{81} could receive a greater proportion of benefits by fighting a battle, so strong type rejects $\hat{x}_t$; this scenario is labeled 2c.

B's decision to accept reveals an important piece of information for A: B is not a strong type. Therefore A_{23} needs to compare her payoff from offering $\overline{x}_t$ (risks war if B_{81} happens to be moderately strong) to the payoff from $\hat{x}_t$ (ensures peace regardless of B_{81}'s type). A_{23} compares the right hand side of the the inequality 2.4 to her belief that $\phi_w = 0.15$: $0.15 < \dfrac{c_A + c_B}{\overline{p}_t - \hat{p}_t + c_A + c_B} = \dfrac{0.01 + 0.015}{0.146 - 0.138 + 0.01 + 0.015} = 0.75$, which means that A_{23}'s belief about B_1's type leads A_{23} to abstain from risking a war. The updated $\pi = 0.153$. A_{23}'s beliefs change, because she has learned B_{81} cannot be strong. Thus, A_{23}'s belief that B_{81} is strong is reduced from 0.15 to 0.05 and the difference is redistributed between A_{23}'s belief that B_{81} is weak, which increases from $\phi_w = 0.15$ to $\phi_w = 0.2$, and the belief that B_{81} is moderately strong increases from $\phi_m = 0.7$ to $\phi_m = 0.75$. The status of the dyad changes to peaceful. In time period 3 and all subsequent rounds, this dyad stays as peaceful, because neither side wants to challenge the new status quo.

Scenario 2b. In those cases when the right hand side of inequality 2.4 is less than A's belief about B being weak, A is willing to risk a war given that her belief about B being weak is relatively large. With respect to the above example, if A_{23}'s belief ϕ_w was large enough to propose $\overline{x}_t$, then peace depends on B_{81}'s type. If B_{81} is weak, then B_{81} accepts and the dyad is characterized as peaceful in the following and subsequent rounds. The new status quo is recorded as $\pi = \overline{x}_t = \overline{p}_t + c_B = 0.146 + 0.015 = 0.161$. A_{23}'s beliefs are updated as $\phi_w = 0.99, \phi_m = 0.05$.

If B_{81} happens to be moderately strong, then he rejects the terms of $\overline{x}_t$, as he would be likely to receive a greater share of the good by fighting a battle. By rejecting, a moderately strong B_{81} signals his type, A_{23}'s beliefs update regardless of the battle outcome: $\phi_w = 0.05, \phi_m = 0.99$. The dyad is then characterized as conflictual and the sides fight a battle in period 3. The battle outcomes are determined probabilistically. If a randomly generated probability $p' > \hat{p}_t$, then A_{23} loses and the new status quo is recorded as $\pi = \hat{x}_t = \hat{p}_t + c_B = 0.138 + 0.015 = 0.153$. Given that B_{81}'s true type is moderately strong, the dyad stays peaceful forever.

However, if $p' \leq \hat{p}_t$ A_{23} happens to win, then π is updated as $\pi = \overline{x}_t = \overline{p}_t + c_B = 0.146 + 0.015 = 0.161$, which means that B_{81} receives fewer benefits from the new status quo than what he potentially could after fighting another battle. Therefore in time period 4, B_{81} will be dissatisfied and the sides will go over the process of negotiation over who gets what again. However, given A_{23}'s updated beliefs about B_{81}'s type, the dyad will not fight again as A_{23} will propose the terms of $\hat{x}_t$.

Scenario 2c. In this scenario, a strong B rejects the terms of $\hat{x}_t$, because he could potentially receive more benefits from fighting a war. Since only a strong type rejects this offer, A's beliefs are updated as $\phi_w = 0.05, \phi_m = 0.05$ and the status of the dyad is set as conflictual in time period 2. In period 3, the dyad fights a battle. With respect to the discussed example, if A_{23} loses, then π is updated as $\pi = \underline{x}_t = \underline{p}_t + c_B = 0.132 + 0.015 = 0.147$, which reflects B_{81}'s type. In time period 4 and all subsequent periods, the dyad stays at peace as neither side could gain more from fighting.

If A_{23} happens to win, then the settlement is recorded as $\pi = \hat{x}_t = \hat{p}_t + c_B = 0.138+0.015 = 0.153$. In period 4, this dyad will be classified as having a conflict of interest, because B_{81}'s type is strong and he could potentially receive more benefits by fighting than π. However, given A_{23}'s updated beliefs, she will propose the terms of $\underline{x}_t$ and therefore the dyad will avoid fighting another armed conflict.

Scenario 3a. Let the beliefs be $\phi_w = 0.7$ and $\phi_m = 0.15$ for the described above example of A_{23} and B_{81}'s dyad. These beliefs generate $\hat{\phi}_w^{k_1} = 0.656$, $\hat{\phi}_w^{k_2} = 0.641$, and $\hat{\phi}_w^{k_3} = 0.269$, which means that $\phi_w > max\{\hat{\phi}_w^{k_2}, \hat{\phi}_w^{k_3}\}$. Given these beliefs, A_{23} proposes terms $\overline{x}_t = \overline{p}_t + c_B = 0.146 + 0.015 = 0.161$ in period 1. This offer is exactly what a weak type of B_{81} would receive from fighting, so if B_{81} happens to be weak, then status quo is updated as $\pi = \overline{x}_t = \overline{p}_t + c_B = 0.161$. In addition, A_{23}'s beliefs are updated to reflect that A_{23} knows B_{81}'s true type: $\phi_w = 0.99, \phi_m = 0.05$. In period 3 and all subsequent time units, the dyad will stay peaceful because it is reflective of B_{81}'s true type and neither side will be dissatisfied with the status quo.

Yet, if B_1 happens to be moderately strong or strong, B_{81} rejects the terms of $\overline{x}_t$, as he could potentially establish a more favorable status quo through fighting. The key feature of these situations is that B_{81}'s rejection does not clearly signal which type B_{81} is, therefore battle outcomes will change A_{23}'s beliefs. These scenarios are discussed below.

Scenario 3b. If B_{81} happens to be moderately strong, B_{81} rejects A_{23}'s terms of $\overline{x}_t$; the dyad is assigned as having an armed conflict. In period 3, if A_{23} wins, then A_{23}'s updated belief about B_{81} being moderately strong is expressed as A_{23}'s probability of victory when B_{81} is moderately strong as a proportion of A_{23}'s total probability of winning: $\phi_m^w =$

$\frac{\hat{p}_t \phi_m}{\hat{p}_t \phi_m + \underline{p}_t(1-\phi_m)} = \frac{0.138 \times 0.15}{0.138 \times 0.15 + 0.132 \times (1-0.15)} \approx 0.16$, which replaces the old belief ϕ_m. If A_{23} loses the battle then the new updated belief about B_1 being moderately strong equals $\phi^l_m = \frac{(1-\hat{p}_t)\phi_m}{(1-\hat{p}_t)\phi_m + (1-\hat{p}_t)(1-\phi_m)} = \frac{0.862 \times 0.15}{0.862 \times 0.15 + 0.868 \times (1-0.15)} = 0.149$. Regardless of the battle outcome, the new belief is compared to the right hand side of inequality 2.5. If the new updated belief is less than or equal to the right hand side of inequality 2.5, then A_{23} believes that B_1 is more likely to be strong and proposes the terms of $\underline{x}_t = 0.147$. A moderately strong type of B_{81} accepts these terms because a moderately strong type is receiving more benefits than otherwise would be the case. The status quo is updated as $\pi = \underline{x}_t = \underline{p}_t + c_B = 0.132 + 0.015 = 0.147$. A's beliefs are updated as $\phi_w = 0.05$ (to reflect that A has found out that B is not weak), and $\phi_m = \phi^w_m$ (to reflect the updated belief ϕ_m after fighting a battle).

If the new belief about B_{81} being moderately strong is greater than the right hand side of the inequality 2.5, then A_{23} believes that B_1 is more likely to be moderately strong and proposes the terms of $\hat{x}_t = 0.153$, which is exactly what a moderately strong B_{81} could expect from fighting. Because A_{23} knows that a strong B_1 would not have accepted such terms, A_{23} updates her beliefs about B_{81}'s type even further, such that $\phi_w = 0.05, \phi_m = 0.99$. The new status quo is recorded as $\pi = \hat{x}_t = \underline{p}_t + c_B = 0.138 + 0.015 = 0.153$, which ensures that the dyad stays peaceful forever as neither side could gain from fighting.

Scenario 3c. If B_{81} happens to be strong, B_{81} rejects A_{23}'s terms of $\overline{x}_t$; the dyad is assigned as having an armed conflict. In period 3, if A_{23} wins, then A_{23}'s updated belief about B_{81} being moderately strong is $\phi^w_m = 0.16$ (see explanation in scenario 3b). If A_{23} loses, the updated belief is $\phi^l_m = 0.149$. If the new updated belief is less than or equal to the right hand side of inequality 2.5, then A_{23} believes that B_{81} is more likely to be strong and proposes the terms of $\underline{x}_t = 0.147$, which are accepted by a strong type of B_{81}. The status quo is updated as $\pi = \underline{x}_t = \underline{p}_t + c_B = 0.132 + 0.015 = 0.147$, the dyad is classified as peaceful. A's beliefs are updated as $\phi_w = 0.05$ (to reflect that A has found out that B is not weak), and $\phi_m = \phi_m^l$ or $\phi_m = \phi_m^w$ (depending on the battle outcome, to reflect the updated belief ϕ_m after fighting a battle)

If the new belief about B_{81} being moderate is greater than the right hand side of the inequality 2.5, then A_{23} believes that B_{81} is more likely to be moderately strong and proposes the terms of $\hat{x}_t = 0.153$, which is less than what a strong B_{81} could expect from fighting and B_{81} rejects the proposal. As A_{23} knows that only a strong type would reject that offer, A_{23}'s beliefs are updated to reflect the reality that only a strong type would have rejected the terms of $\hat{x}_t$: $\phi_w = 0.05, \phi_m = 0.05$. Still, battles are probabilistic, which means that if B_1 wins, then the status quo is updated as $\pi = \underline{x}_t = \underline{p}_t + c_B = 0.132 + 0.015 = 0.147$, which ensures peace forever. Yet, if B_{81} loses then the status quo is updated as $\pi = \hat{x}_t = \hat{p}_t + c_B = 0.138 + 0.015 = 0.153$, which means that in period 5, the dyad will be classified as having a conflict of interest, because B_{81}'s strong type determines that B_1 will be dissatisfied with the agreement π. However, given A_{23}'s new updated beliefs $\phi_w = 0.05, \phi_m = 0.05$, it is highly likely that the conflict will be resolved without their resort to arms.

Table 2.2 summarizes the key features and quantities of interest of the actors in the benchmark model as discussed above. I group the variables as describing either the world, a dyad, or an individual agent.

2.3.3 Analysis of the exhaustive parameter space

In the discussion of the illustrative runs, I focused on the influence of the initial beliefs of side A and the process of their updating. Now, I turn to analyzing the model systematically, i.e., focusing on the effects of major parameters across the entire range of their values. As the benchmark model can be solved analytically, I analyze numerical output of the model to *verify* the computer program by demonstrating that the computer code accomplishes what the model assumptions imply. I also use the analysis of the numerical output from the null model to verify the genetic automated search program introduced in chapter 3 by comparing the results from the exhaustive parameter space and the partial parameter space.

To demonstrate that the computer code implements the sequences of peaceful, potentially conflictual, and conflictual interactions described in this chapter, I focus on (1) the effect of the costs that the sides have to pay to fight a battle on the dyad's probability of entering a potential conflict or an armed conflict and (2) the effect of the distribution of capabilities in the dyad on the probability of entering a potentially conflictual phase and conflictual phase. The model implies that (holding the beliefs constant) A is more likely to propose the terms that risk an armed conflict as fighting becomes cheaper and as the difference between the probability of A defeating two different types of B increases. These two major relationships are implied in A's decision-making rules (expressed as inequalities 2.1-2.5), so the analysis that follows does not discover these relationships, rather it demonstrates that the code is implemented correctly.

The null model's parameters are summarized in Table 2.1. A parameter is an attribute of a model whose initial value is used to create the actual model. It describes an initial model condition (*How to Use Parameters and Parameter files. RePast Documentation* N.d.). For example, the parameter ϕ_w specifies the initial belief of A, which may change over the course of the simulation model's execution. In contrast, N_a sets the number of agents in the system, and this condition is preserved throughout the "life" of the system. The exhaustive search space of this model is fairly large: the effects of nine parameters are explored throughout their respective ranges of values. With respect to the costs that the belligerents pay to fight a battle, parameters c_A and c_B express the costs for side A and side B of any given dyad. The costs of fighting a single battle cannot exceed 30% of the belligerent's observable capabilities. Both c_A and c_B were varied across their ranges of $[0, 0.1]$ with an increment of 0.01, with the minimum value of 0.01 and the maximum of 0.1. With respect to the difference in A's probability of victory against two different types of B, five parameters account for that: A's total capabilities m_A, B's observable capabilities m_B, and also $\underline{b}, \hat{b}, \overline{b}$ which record B's private capabilities with regard to B's type. The ranges of A's total and B's observable capabilities are both $[0, 1.0]$; these parameters are varied across this range with an increment of 0.1 such that the minimum value is 0.1 and the maximum is 1.0. The ranges of B's private capabilities are all $[0, 0.1]$; in addition, these values should each be less than B's observable capabilities and be generated such that $\underline{b} < \hat{b} < \overline{b}$. These values were varied across their ranges with an increment of 0.01, such that the minimum value is 0.01 and the maximum is 0.1. Given the restrictions imposed on the size of c_A, c_B, $\underline{b}, \hat{b}$, and $\overline{b}$, Tables 2.3 and 2.4 show the possible ranges of these variable given the mentioned restrictions and a sample parameter file that accounts for these restrictions.

Had there been no restrictions on the parameter values, then nine variables with ten possible values each would require $10^9 = 1,000,000,000$ unique combinations of initialization parameter values. However, given all the restrictions mentioned above, Table 2.4 specifies that given the chosen increments the nine parameters need almost 18 million unique runs to create the full parameter space. However, each simulation run is unique given the randomized values of certain variables (e.g., the decision which side wins a battle is executed by generating a random value and comparing it to A's probability of winning). Therefore, I run each unique combination of parameter values at least 10 times to ensure that the analyzed trends are representative of the relationships in the model, rather than being an extreme case. So, the total number of runs executed is $17,968,500 \times 10 = 179,685,000$. After having created the full search space, I collapse the mean values by their unique combinations of parameter values. The analyzed data set has ≈18 million observations, where each observation represents the mean outcome of ten simulations. Thus, each observation includes the values of the unique combination of initialization parameter values and the corresponding average value of the response variable that the simulation produces given the initialization parameter values.

To assess the main relationships I construct three variables. *Parity* measures how equally matched are the two sides. I code $Parity_t = \frac{m_{Ht}}{m_{Ht} + m_{Lt}}$, where $m_{Ht} = m_A$, $m_{Lt} = m_B + b$ if $m_A > (m_B + b)$ and $m_{Ht} = m_B + b$, $m_{Lt} = m_A$ if $m_A < (m_B + b)$. I use the value of b with respect to B's type. The variable ranges from 0 to 0.5. The model implies that greater values of parity lead to higher probability of A proposing the terms that risk fighting an armed conflict.

Difference in A's probability of defeating two different types of B is related to *Parity*. Recall that I assumed that A's probability of winning a battle equals the proportion of the dyad's total capabilities that A possesses. If A and B's capabilities are similar, then the difference between A's probabilities of winning against two different types is greater than when A and B's capabilities are different. *Difference*$_{\overline{p}_{t-1},\hat{p}_{t-1}}$ equals $\overline{p}_{t-1} - \hat{p}_{t-1}$. *Difference*$_{\hat{p}_{t-1},\underline{p}_{t-1}}$ is $\hat{p}_{t-1} - \underline{p}_{t-1}$. *Difference*$_{\overline{p}_{t-1},\underline{p}_{t-1}}$ $= \overline{p}_{t-1} - \underline{p}_{t-1}$. The model implies that higher values of *Difference* increase A's probability of risking war, and thus dyad's fighting a war.

Sum of costs is expressed as the sum of c_A and c_B. The model assumes that the more expensive it is for the belligerents to fight, the less likely they are to fight. In contrast, the cheaper the battle, the more likely A is to propose the terms that risk war.

The logistic regressions that estimate the effects of these covariates on dyad's fighting an armed conflict in time t are shown in Table 2.5. The three models differ by the choice of variable *Difference*, which either describes the difference between A' probability to win against a strong type and a moderately strong type, or the difference between A' probability to win against a moderately strong type and a weak B, or the difference between A' probability of military victory against a strong type and a weak type. The results validate that the computer code is implemented as described in this chapter. *Parity*$_{t-1}$ increases the probability of a dyad to fight a battle in time t. The substantive effect of this variable is great. In model 1, the logit estimate of $\hat{\beta}$ is 8.24. A ten point change in *Parity*$_{t-1}$ (expressed as a percentage) in the odds that a dyad fights a battle equals $100[exp(\hat{\beta}\delta) - 1]$=6.1075258e+37 percent. All types of *Difference* generate a similarly huge substantive effect on the probability of a dyad fighting an armed conflict, which is exactly what the data are supposed to show given that these two measures assess the same concept. In model 3, the logit estimate of $\hat{\beta}$ for *Difference* is 3.62. A five point change in *Difference*$_{t-1}$ (expressed as a percentage) in the odds that a dyad fights a battle equals 7,256,548,737.23 percent. The *Sum of costs* decreases the probability of a dyad fighting an armed conflict. Again, the effect is huge. For instance, in model 2, the logit estimate of $\hat{\beta}$ for *Sum of costs* is -162.12. A ten unit change in *Sum of costs*$_{t-1}$ (expressed as a percentage) in the odds that a dyad fights a battle equals $100[exp(\hat{\beta}\delta) - 1] \approx -100\%$. It is important to emphasize that these effects are not inferential in the traditional sense, but merely describe the relationships in the model, that are assumed in the decision-making rules that may lead to war. This analysis validates my implementation of the model in the Java programming language.

2.4 Conclusion

In this chapter, I have laid out the theoretical foundations of the benchmark model that simulates peaceful, potentially conflictual, and conflictual interactions between actors in an artificial international system. The model assumes uncertainty that side A of each dyad has over B's private component of capabilities. All else equal, (1) when the sides' capabilities are close to parity (and by implication, the greater is A's probability of defeating a weaker type of opponent relative to the stronger type), (2) when the costs of fighting are low, (3) when A's belief that B is weak or moderately strong are relatively large compared to the belief that B is strong, A is more likely propose the terms that risk war if rejected by a relatively stronger type of opponent. I have laid out the pseudocode of the computer simulation, explained the effects of A's beliefs through the illustrative runs, and summarized the relationships between the cost of fighting and parity of capabilities on the one side and the probability of violent conflict on the other side by analyzing the exhaustive parameter space of the benchmark computational model.

Fig. 2.1. Counts of peaceful dyads, dyads that experience a conflict of interests, and dyads at armed conflict shown over time in a single simulation run

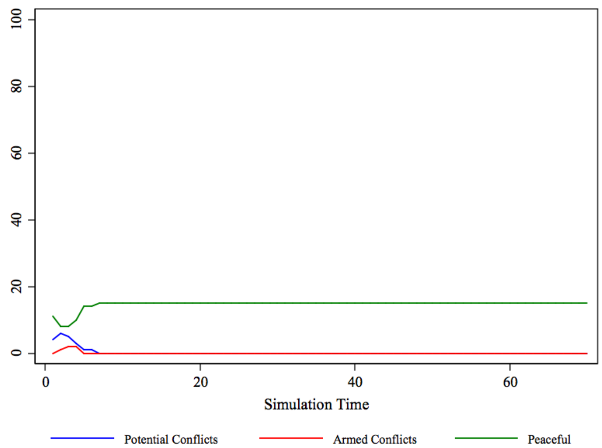

(a) “World” system of 15 dyads

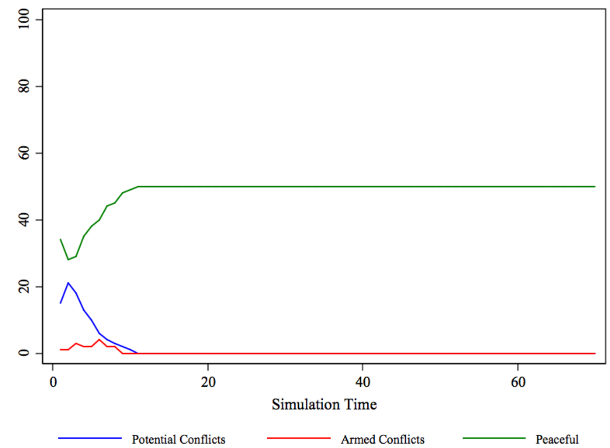

(b) “World” system of 50 dyads

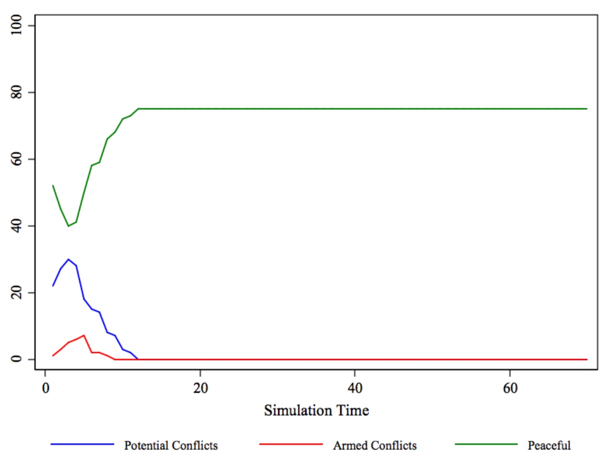

(c) “World” system of 75 dyads

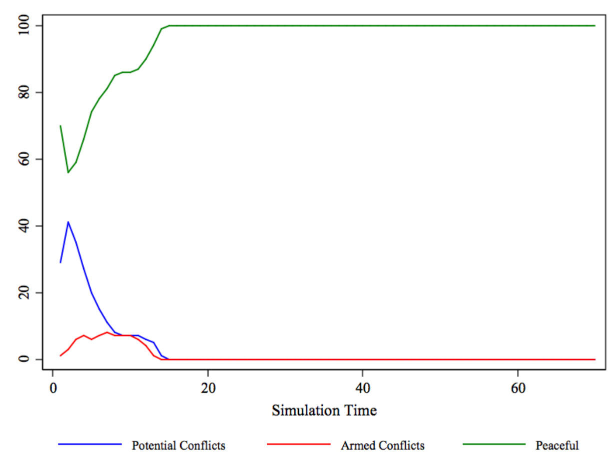

(d) “World” system of 100 dyads

Note: regardless of the number of dyads in the artificial international system the major trend remains: initially, some proportion of dyads experiences conflicts of interests, and some of those disputes become violent. After the initial 10-17 time units, the conflicts of interests no longer arise, as they get settled through redistributing the benefits in accordance with the distribution of relative power.
Analogous time series are shown for the extended model in Figure 4.1.

Fig. 2.2. Change in A's beliefs over time, shown for 50 dyads, color coded by dyad

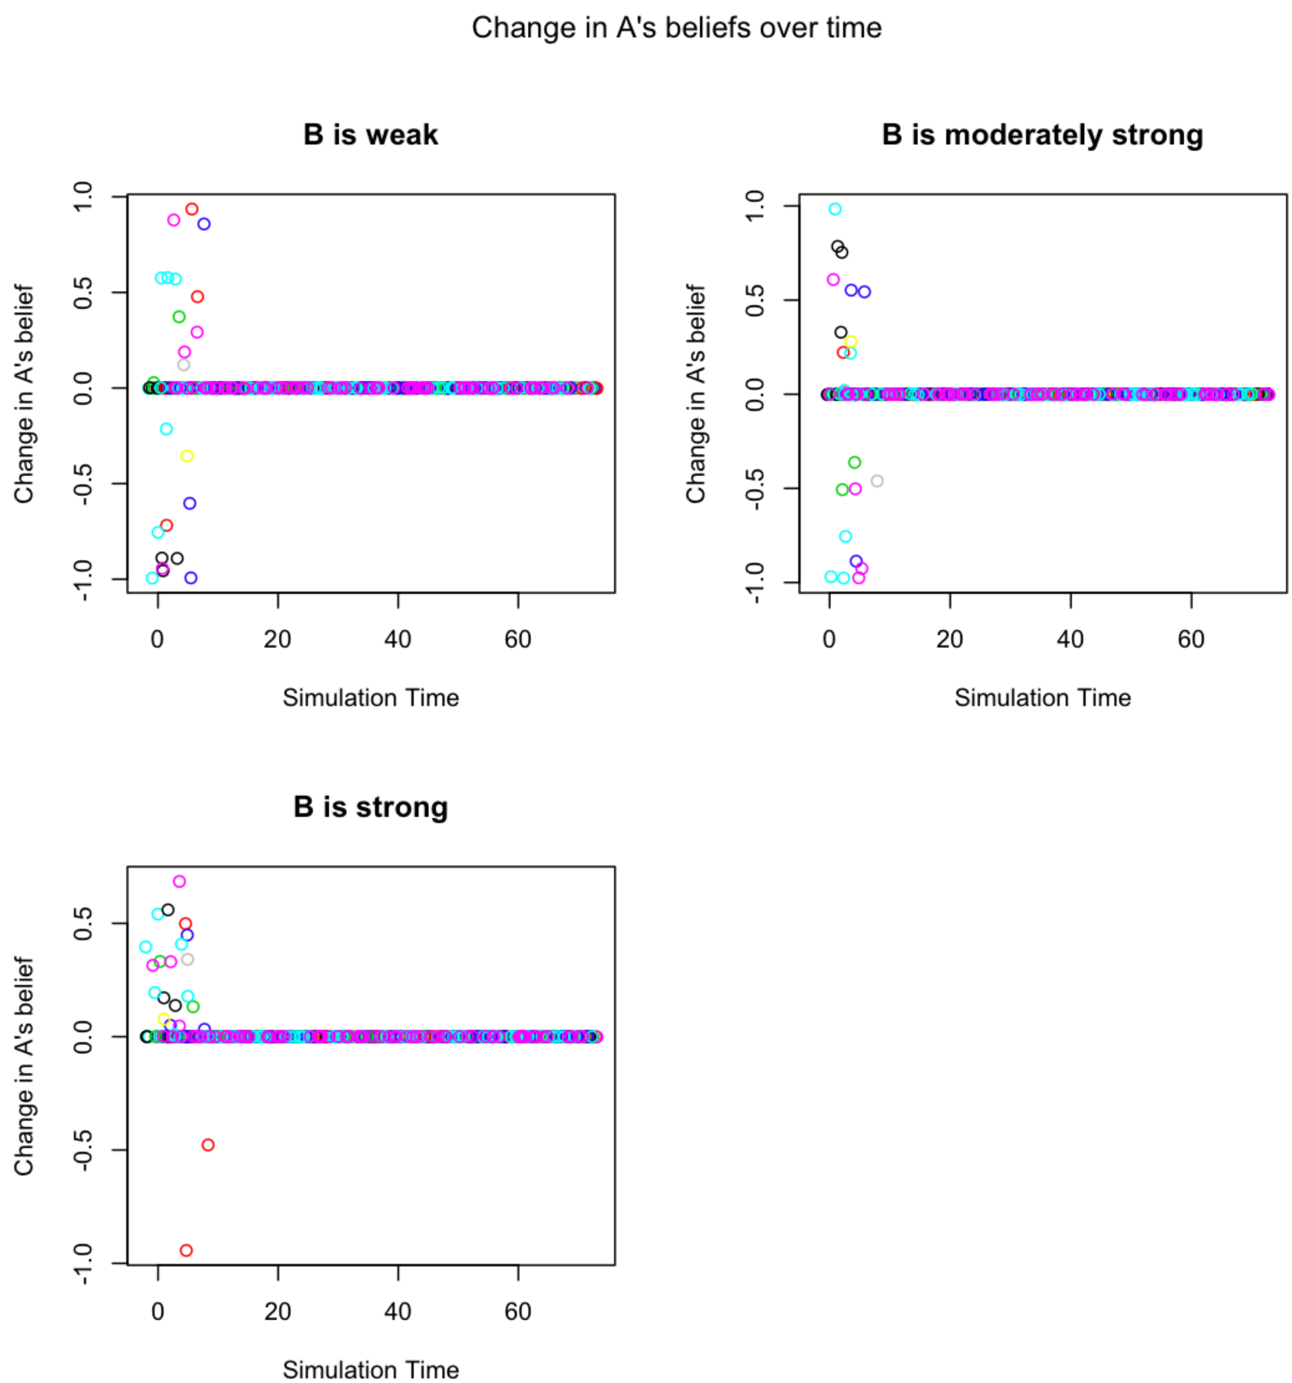

Note: the graphs chart the differences between A's beliefs about B's strength in time period t from A's beliefs in time period $t-1$ (i.e., $\phi_t - \phi_{t-1}$). After the initial 10-17 time units, the "learning" of B's type is complete. It is important to note that A's belief that B is strong rarely decreases, as wars break out due to A's underestimation of B's strength. The only scenario that allows for decrease in A's belief that B is not as strong as previously thought is described in step **??** of the pseudocode.

Fig. 2.3. Absolute change in A's beliefs over time, shown for 50 dyads

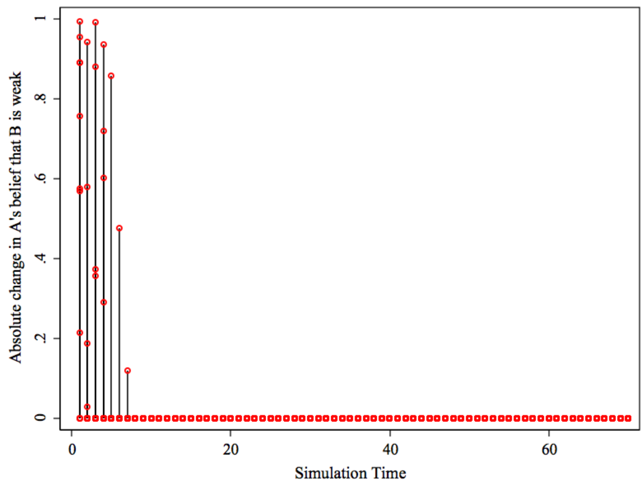

(a) Absolute change in A's belief that B is weak

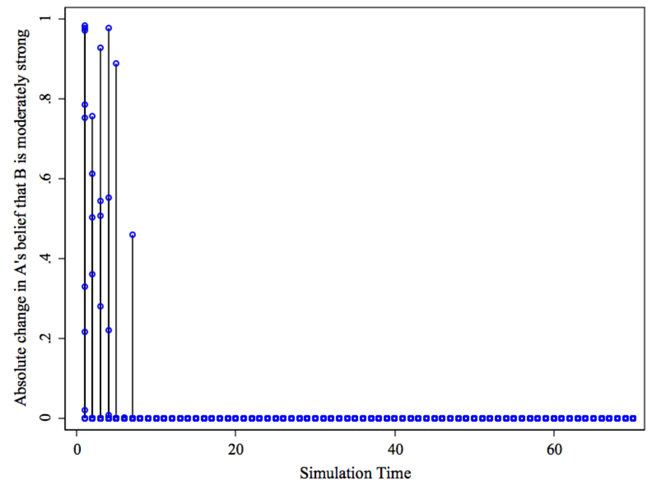

(b) Absolute change in A's belief that B is moderately strong

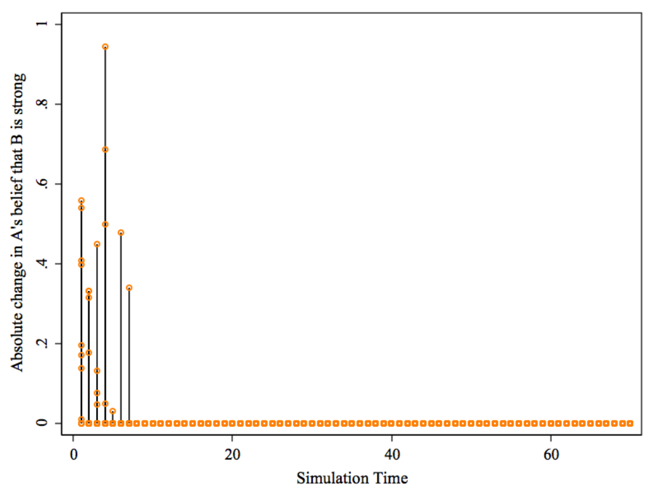

(c) Absolute change in A's belief that B is strong

Note: the graphs chart absolute differences between A's beliefs about B's strength in time period t from A's beliefs in time period $t-1$, i.e., $\phi_t - \phi_{t-1}$. For comparison, consider absolute changes in A's beliefs in the extended model in Figure 4.2.

Fig. 2.4. Changes in the division of the disputed good, π, over time, shown for 50 dyads

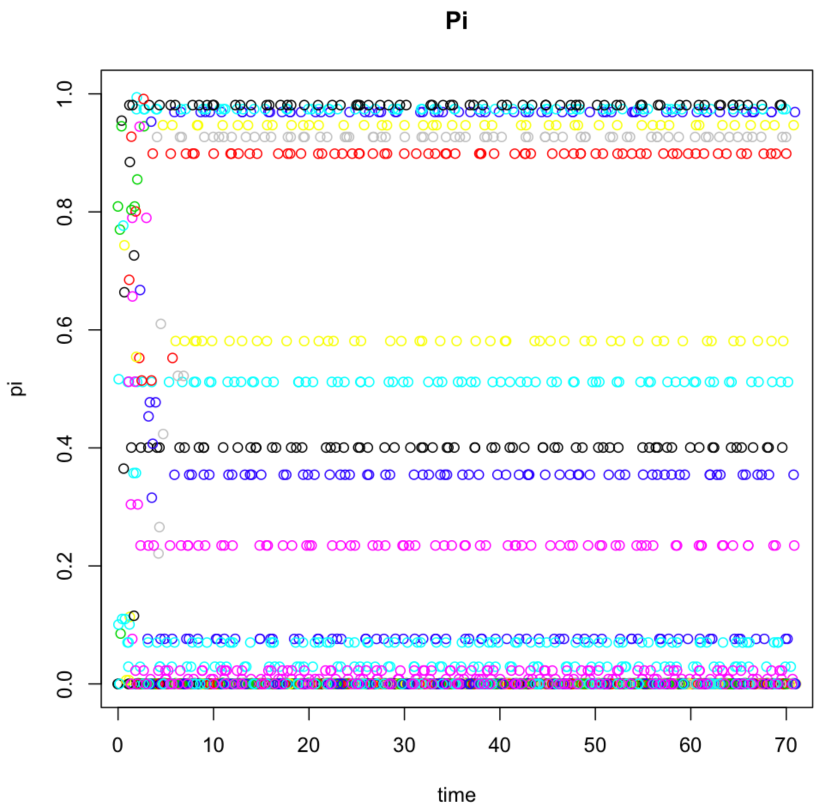

(a) Division of the disputed good π shown over time, color coded by dyad

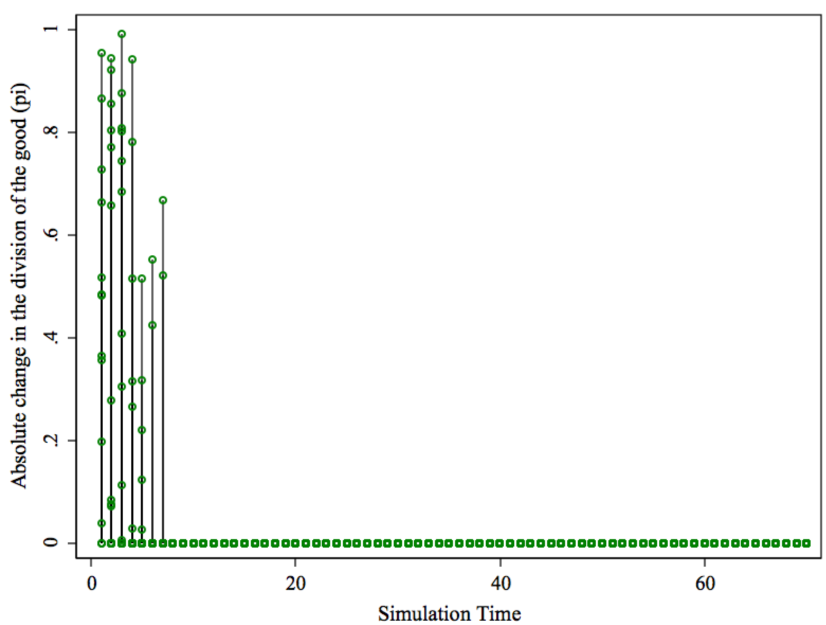

(b) Absolute change in the division of the disputed good between t and $t-1$, i.e., $\pi_t - \pi_{t-1}$

Note: For comparison of how π changes in the extended model, see Figure 4.3.

Fig. 2.5. The comparison of the simulated capabilities to the empirical distribution

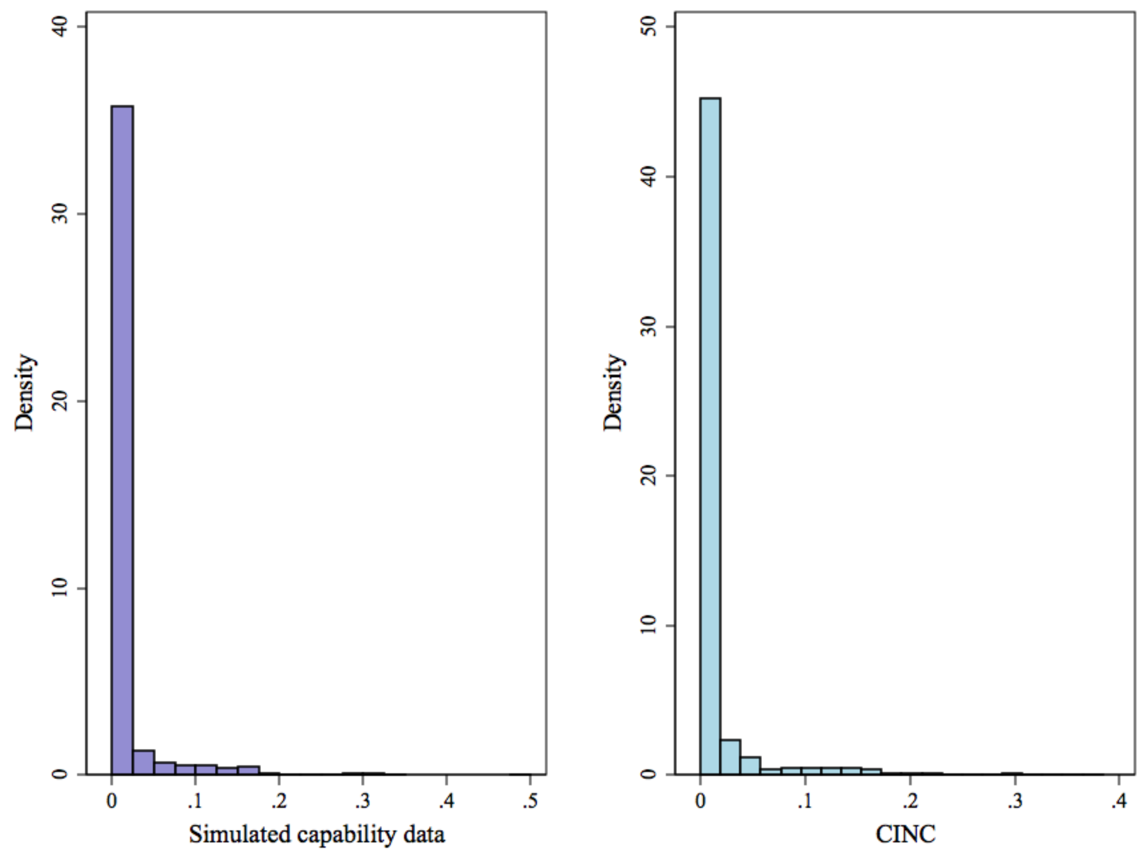

Note: distribution of simulated capability vs. empirical distribution of Composite Index of National Capability (CINC, Singer, Bremer and Stuckey (1972).

Table 2.1. Initialization parameters in the null model

Parameter	Code Reference	Range	Default Value	Interpretation
N_a	`num agents`	$[2, \infty]$	200	Number of actors
N_d	`num dyads`	$[1, \infty]$	50	Number of dyads
π	`pi`	$[0, 1]$ SD=0.288	Uniformly distributed,	A's share of benefits
m_A	`caps A`	$[0, 1]$	Beta distributed, $\alpha = 0.05$, $\beta = 1.25$	A's total capabilities
m_B	`caps B`	$[0, 1]$	Beta distributed, $\alpha = 0.05$, $\beta = 1.25$	B's observable capabilities
${\phi_w}^{\dagger}$	`phi weak`	$[0, 1]$	Uniformly distributed, SD=0.288	A's belief that B is weak
${\phi_m}^{\dagger}$	`phi mod`	$[0, 1]$	Uniformly distributed, SD=0.288	A's belief that B is strong
$\underline{b}^{\S}$	`underline b`	$[0, 0.1]$	Uniformly distributed, SD=0.029	B's private capabilities $\underline{b} < m_B$
$\hat{b}^{\S}$	`hat b`	$[0, 0.1]$	Uniformly distributed, SD= 0.029	B's private capabilities $\hat{b} < m_B$
$\overline{b}^{\S}$	`overline b`	$[0, 0.1]$	Uniformly distributed, SD= 0.029	B's private capabilities $\overline{b} < m_B$
c_A	`cost A`	$[0, 0.1]$	Uniformly distributed, SD= 0.029	A's cost of fighting a battle, $c_A < 0.3 \times m_A$
c_B	`cost A`	$[0, 0.1]$	Uniformly distributed, SD= 0.029	B's cost of fighting a battle, $c_B < 0.3 \times m_B$

Note: These are the key parameters that may be systematically varied across simulations;
† indicates that the following condition applies: $\phi_w + \phi_m \leq 1$;
§ indicates that the following conidtion applies: $[(\underline{b} < \hat{b} < \overline{b})$ AND $(\underline{b}||\hat{b}||\overline{b}) < 30\%$ of $m_B]$

These parameters are all present in the extended version of the model as well, see Table 4.1 for additional initialization parameters introduced in the extended model.

Table 2.2. Characteristics and quantities of interest that describe the world, dyads, and actors in the null model

Characteristic of the world	
N_a	Number of actors
N_d	Number of dyads
Characteristics of dyads	
π	The proportion of the disputed good that A controls. Exogenously given.
ϕ_w	A's belief that B is weak. Exogenously given.
ϕ_m	A's belief that B is moderately strong. Exogenously given.
$\underline{x}_t$	The terms proposed by A in period t if A believes that B is strong
$\hat{x}_t$	The terms proposed by A in period t if A believes that B is moderately strong
$\overline{x}_t$	The terms proposed by A in period t if A believes that B is weak
$\hat{\phi}_w^{k_1}$	The cutpoint 1 that determines whether A prefers $\underline{x}_t$ to $\hat{x}_t$
$\hat{\phi}_w^{k_2}$	The cutpoint 2 that determines whether A prefers $\underline{x}_t$ to $\overline{x}_t$
$\hat{\phi}_w^{k_3}$	The cutpoint 3 that determines whether A prefers $\hat{x}_t$ to $\overline{x}_t$
$\phi_m^{\prime}$	A's updated belief that B is moderately strong.
Characteristics of actors	
m_A	A's total capabilities.
m_B	B's publicly observable capabilities.
$\underline{b}$	B's private capabilities, if B is a weak type. Drawn from a uniform or normal distribution, such that $\underline{b} < \hat{b} < \overline{b}$.
$\hat{b}$	B's private capabilities, if B is a moderately strong type. Drawn from a uniform or normal distribution, such that $\underline{b} < \hat{b} < \overline{b}$.
$\overline{b}$	B's private capabilities, if B is a strong type. Drawn from a uniform or normal distribution, such that $\underline{b} < \hat{b} < \overline{b}$.
$\overline{p}_t$	$= \dfrac{m_A}{m_A + m_B + \underline{b}}$, the probability of A's winning a battle if B is a weak type.
$\hat{p}_t$	$= \dfrac{m_A}{m_A + m_B + \hat{b}}$, the probability of A's winning a battle if B is moderately strong
$\underline{p}_t$	$= \dfrac{m_A}{m_A + m_B + \overline{b}}$, the probability of A's winning a battle if B is strong.
c_A	A's cost of fighting a battle
c_B	B's cost of fighting a battle

Notes: These characteristics and quantities of interest are all present in the extended version of the model that incorporates the influence of third parties on a given dyad; see Table 4.2 for additional characteristics and quantities of interest introduced in the extended model.

Table 2.3. Initialization parameters varied in the model

Parameter	Range	Discrete values used	Num possible values
m_A	$(0,1]$	0.1,0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0	10
c_A	$(0,0.1]$	0.01,0.02, 0.03, 0.04, 0.05, 0.06, 0.07, 0.08, 0.09, 0.1	10
		Condition: $c_A < 0.3 \times m_A$	
		IF $m_A = 0.1$, c_A is varied 0.01 to 0.03	3
		IF $m_A = 0.2$, c_A is varied 0.01 to 0.06	6
		IF $m_A = 0.3$, c_A is varied 0.01 to 0.09	9
m_B	$(0,1]$	0.1,0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0	10
c_B	$(0,0.1]$	0.01,0.02, 0.03, 0.04, 0.05, 0.06, 0.07, 0.08, 0.09, 0.1	10
		Condition: $c_B <$ 30% of m_B	
		IF $m_B = 0.1$, c_B is varied 0.01 to 0.03	3
		IF $m_B = 0.2$, c_B is varied 0.01 to 0.06	6
		IF $m_B = 0.3$, c_B is varied 0.01 to 0.09	9
$\underline{b}$*	$(0,0.1]$	0.01,0.02, 0.03, 0.04, 0.05, 0.06, 0.07, 0.08	8
$\hat{b}$*	$(0,0.1]$	0.01,0.02, 0.03, 0.04, 0.05, 0.06, 0.07, 0.08, 0.09	9
$\overline{b}$*	$(0,0.1]$	0.01,0.02, 0.03, 0.04, 0.05, 0.06, 0.07, 0.08, 0.09, 0.1	10
		Conditions: $\underline{b} < \hat{b} < \overline{b}$ & $b <$ 30% of m_B	
		IF $m_B = 0.1$, $\overline{b}$ is varied 0.01 to 0.09	9
		$\hat{b}$ is varied 0.01 to 0.08	8
		$\underline{b}$ is varied 0.01 to 0.07	7
ϕ_w*	$(0,0.99]$	0.005, 0.1,0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.99	11
ϕ_m*	$(0,0.99]$	0.005, 0.1,0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9,0.99	11
		Condition: $\phi_w + \phi_m \leq 1$	
		IF $\phi_w = 0.005$, ϕ_m is varied 0.05 to 0.99	11
		IF $\phi_w = 0.1$, ϕ_m is varied 0.05 to 0.9	10
		IF $\phi_w = 0.2$, ϕ_m is varied 0.05 to 0.8	9
		IF $\phi_w = 0.3$, ϕ_m is varied 0.05 to 0.7	8
		IF $\phi_w = 0.4$, ϕ_m is varied 0.05 to 0.6	7
		IF $\phi_w = 0.5$, ϕ_m is varied 0.05 to 0.5	6
		IF $\phi_w = 0.6$, ϕ_m is varied 0.05 to 0.4	5
		IF $\phi_w = 0.7$, ϕ_m is varied 0.05 to 0.3	4
		IF $\phi_w = 0.8$, ϕ_m is varied 0.05 to 0.2	3
		IF $\phi_w = 0.9$, ϕ_m is varied 0.05 to 0.1	2

Table 2.4. Sample parameter file to execute exhaustive search of the null model

Parameter	Range	Discrete values used	Number of combinations
m_A	$(0, 0.4]$	0.0001, 0.001, 0.05, 0.1, 0.15, 0.2, 0.25. 0.3, 0.35, 0.4	
c_A	$[0.00003, 0.12]$	IF $-_A$= 0.0001, c_A =0.00003	1
		IF m_A= 0.001, c_A =0.00003,0.0003	2
		IF m_A= 0.05, c_A =0.00003, 0.0003, 0.0015	3
		IF m_A= 0.1, c_A =0.00003, 0.0003, 0.0015, 0.03	4
		IF m_A= 0.15, c_A =0.00003, 0.0003, 0.0015, 0.03, 0.045	5
		IF m_A= 0.2, c_A =0.00003, 0.0003, 0.0015, 0.03, 0.045 0.06	6
		IF m_A= 0.25, c_A =0.00003, 0.0003, 0.0015, 0.03, 0.045 0.06, 0.075	7
		IF m_A= 0.25, c_A =0.00003, 0.0003, 0.0015, 0.03, 0.045 0.06, 0.075, 0.09	8
		IF m_A= 0.3, c_A =0.00003, 0.0003, 0.0015, 0.03, 0.045 0.06, 0.075, 0.09, 0.105	9
		IF m_A= 0.3, c_A =0.00003, 0.0003, 0.0015, 0.03, 0.045 0.06, 0.075, 0.09, 0.105, 0.12	10
		Sum of unique combinations =	55
m_B	$(0, 0.4]$	0.0001, 0.001, 0.05, 0.1, 0.15, 0.2, 0.25. 0.3, 0.35, 0.4	
c_B	$[0.00003, 0.12]$	IF m_B= 0.0001, c_B =0.00003	1
		IF m_B= 0.001, c_B =0.00003,0.0003	2
		IF m_B= 0.05, c_B =0.00003, 0.0003, 0.0015	3
		IF m_B= 0.1, c_B =0.00003, 0.0003, 0.0015, 0.03	4
		IF m_B= 0.15, c_B =0.00003, 0.0003, 0.0015, 0.03, 0.045	5
		IF m_B= 0.2, c_B =0.00003, 0.0003, 0.0015, 0.03, 0.045 0.06	6
		IF m_A= 0.25, c_A =0.00003, 0.0003, 0.0015, 0.03, 0.045 0.06, 0.075	7
		IF m_A= 0.25, c_A =0.00003, 0.0003, 0.0015, 0.03, 0.045 0.06, 0.075, 0.09	8
		IF m_A= 0.3, c_A =0.00003, 0.0003, 0.0015, 0.03, 0.045 0.06, 0.075, 0.09, 0.105	9
		IF m_A= 0.3, c_A =0.00003, 0.0003, 0.0015, 0.03, 0.045 0.06, 0.075, 0.09, 0.105, 0.12	10
		Sum of unique combinations =	55
$\overline{b}$		IF m_B = 0.0001, $\overline{b}$=0.00003,0.000015,0.00001	3
$\hat{b}$		$\hat{b}$=0.000015, 0.00001, 0.000005	3
$\underline{b}$		$\underline{b}$= 0.000003, 0.000001, 0.0000005	3
$\overline{b}$		IF m_B = 0.001, $\overline{b}$=0.0003 ,0.00015,0.00001	3
$\hat{b}$		$\hat{b}$=0.00015, 0.0001, 0.00001	3
$\underline{b}$		$\underline{b}$= 0.00003, 0.00001, 0.000005	3

Continued on next page

Continued from previous page			
Parameter	Range	Discrete values used	Number of combinations
$\overline{b}$		IF $m_B = 0.005$, $\overline{b}$= 0.0015, 0.0005, 0.0001	3
$\hat{b}$		$\hat{b}$=0.00015, 0.00005, 0.00001	3
$\underline{b}$		$\underline{b}$=0.00003, 0.000005, 0.000005	3
$\overline{b}$		IF $m_B = 0.1$, $\overline{b}$=0.03,0.015,0.005	3
$\hat{b}$		$\hat{b}$=0.015,0.001,0.0005	3
$\underline{b}$		$\underline{b}$= 0.001,0.0005,0.00005	3
$\overline{b}$		IF $m_B = 0.15$, $\overline{b}$= 0.045 , 0.023 , 0.001	3
$\hat{b}$		$\hat{b}$= 0.023 , 0.003 , 0.0005	3
$\underline{b}$		$\underline{b}$=0.003 , 0.0005 , 0.0001	3
$\overline{b}$		IF $m_B = 0.2$, $\overline{b}$= 0.06, 0.03, 0.015	3
$\hat{b}$		$\hat{b}$=0.03, 0.015, 0.001	3
$\underline{b}$		$\underline{b}$= 0.015, 0.001, 0.0005	3
$\overline{b}$		IF $m_B = 0.25$, $\overline{b}$= 0.075, 0.05, 0.025	3
$\hat{b}$		$\hat{b}$= 0.05, 0.025, 0.015	3
$\underline{b}$		$\underline{b}$= 0.025, 0.015, 0.005	3
$\overline{b}$		IF $m_B = 0.3$, $\overline{b}$= 0.09, 0.045, 0.025	3
$\hat{b}$		$\hat{b}$= 0.045, 0.025, 0.01	3
$\underline{b}$		$\underline{b}$= 0.025, 0.01, 0.05	3
$\overline{b}$		IF $m_B = 0.35$, $\overline{b}$=0.105,0.05,0.025	3
$\hat{b}$		$\hat{b}$=0.05,0.025, 0.01	3
$\underline{b}$		$\underline{b}$= 0.025, 0.01, 0.005	3
$\overline{b}$		IF $m_B = 0.4$, $\overline{b}$=0.12, 0.06, 0.03	3
$\hat{b}$		$\hat{b}$=0.06, 0.03, 0.015	3
$\underline{b}$		$\underline{b}$=0.03, 0.015, 0.01	3
		Sum of unique combinations =	90
ϕ_w	$(0, 0.99]$	IF $\phi_w = 0.005$, ϕ_m is varied 0.05 to 0.99	11
ϕ_m	$(0, 0.9]$	IF $\phi_w = 0.1$, ϕ_m is varied 0.05 to 0.9	10
ϕ_m	$(0, 0.8]$	IF $\phi_w = 0.2$, ϕ_m is varied 0.05 to 0.8	9
ϕ_m	$(0, 0.7]$	IF $\phi_w = 0.3$, ϕ_m is varied 0.05 to 0.7	8
ϕ_m	$(0, 0.6]$	IF $\phi_w = 0.4$, ϕ_m is varied 0.05 to 0.6	7
ϕ_m	$(0, 0.5]$	IF $\phi_w = 0.5$, ϕ_m is varied 0.05 to 0.5	6
ϕ_m	$(0, 0.4]$	IF $\phi_w = 0.6$, ϕ_m is varied 0.05 to 0.4	5
ϕ_m	$(0, 0.3]$	IF $\phi_w = 0.7$, ϕ_m is varied 0.05 to 0.3	4
ϕ_m	$(0, 0.2]$	IF $\phi_w = 0.8$, ϕ_m is varied 0.05 to 0.2	3
ϕ_m	$(0, 0.1]$	IF $\phi_w = 0.9$, ϕ_m is varied 0.05 to 0.1	2
ϕ_m	$(0, 0.05]$	IF $\phi_w = 0.99$, ϕ_m is 0.05	1
		Sum of unique combinations =	66
		Total of unique runs (the product of sums)=	17,968,500

Note: See Table 4.4 for the discrete values of the initialization parameters unique for the extended model.

Table 2.5. Logistic regressions of armed conflict onset

	1			**2**			**3**		
Covariate	β	(S.E.)	p-value	β	(S.E.)	p-value	β	(S.E.)	p-value
$Parity_{t-1}$	9.27	(.87)	0.000	9.01	(.93)	0.000	8.06	(.89)	0.000
$Difference_{\overline{p}_{t-1},\hat{p}_{t-1}}$	5.44	(.53)	0.000						
$Difference_{\hat{p}_{t-1},\underline{p}_{t-1}}$				6.28	(.57)	0.000			
$Difference_{\overline{p}_{t-1},\underline{p}_{t-1}}$							3.62	(.42)	0.000
$Sum\ of\ costs_{t-1}$	-209.67	(63.31)	0.000	-162.12	(60.02)	0.000	-187.31	(67.58)	0.000
Constant	-3.27	(.14)	0.000	-3.43	(.17)	0.000	-3.32	(.17)	0.000
Log likelihood		-371.27			-344.18			-350.66	
LR χ^2		150.92			205.12			192.15	
N					17,968,500				

Note: The number of observations does not match the number of unique simulation runs, because of some unique combinations were lost in data transfer. Each observation in the data represents the mean outcome of ten simulations. Due to the size of data, the processing and computation were done remotely via the Penn State HPC system.

Chapter 3

Genetic Automated Search for Complex Simulation Models

In the previous chapter I have introduced the benchmark agent-based model of conflict management and presented the analysis of its exhaustive search space. This chapter introduces a software package designed to address a large parameter space problem that is at the root of the difficulties that analysts experience when interpreting complex simulation models in a systematic fashion. Systematic interpretation of complex simulation models has been done through generating exhaustive parameter spaces, which are often prohibitively costly to create. By omitting extensive robustness checks to sets of parameters, or leaving out the analysis of extreme conditions under which a model "breaks down," an analyst risks presenting incomplete analyses of his or her model. The customizable software program that I develop interfaces a genetic algorithm search with simulation models, treating any simulation model as a fitness function of the genetic algorithm. Genetic algorithms do not need to create an exhaustive parameter space to find global solutions, and therefore users are able to explore their computational models in less time, yet systematically.

I first outline the difficulties of systematic interpretation of computational models, explain the basics of the genetic algorithm as an optimization heuristic (section 3.2) and then proceed to outlining the structure and the logic of the *SimGA* software program, which offers a general solution to model validation as the program may be coupled with any external simulation model that may be executed in a batch mode and initialized with a parameter file (section 3.3.1). An application to the benchmark model follows in section 3.3.2.

The genetic search approach is more efficient than creating an exhaustive parameter space, however just like the exhaustive parameter space needs further interpretation with statistical methods, the output of solutions from the genetic search program requires further exploration as well. In section 3.3.2, I describe some of the simple statistical approaches to discern the patterns within the solutions created by the genetic algorithm program: superimposed kernel density plots and correlation matrices. One of the necessary additions for this software application is a pattern recognition module, which is especially important for more complex models, as discussed in chapter 4. I discuss this addition further in subsection 6.3.2.

3.1 A Large Parameter Space Problem

Complex simulation models do not have analytical solutions and require analysis of numerical output from multiple simulation runs.[1] Traditionally, scholars have been generating however many possible combinations of parameter values, ideally creating an exhaustive parameter space, and then searching through such enormous sets of data using statistical techniques and comparative statics. Due to their size, creating exhaustive search spaces is costly, both in terms of time and computational power. Yet, exploring the full space is critical for uncovering the "extreme" predictions of the model, and exploring the model's robustness to groups of parameters. In computational political science, there is a need for

[1] Hereafter, I use the terms "complex simulation model" or "simulation model" as a superset of any computational model that represents some aspect of social life. In Political Science, the most common types of complex simulation models is agent-based or multi-agent models that are characterized by the presence of many individual agents and emergent dynamics (Miller and Page 2007).

a method that would enable analysts to explore their simulations for extreme predictions, robustness, and nonlinear interactions of variables in a more efficient fashion than creating an exhaustive parameter space.

Large parameter spaces are time consuming to generate. To generate an exhaustive parameter space, an analyst has to iterate over the initialization parameters with small enough increments to capture the changes in the response variable that may result from small fluctuations in the parameter values and yet with large enough increments for the task to be manageable. In addition, simulation models are non-deterministic due to randomization steps within models and require multiple iterations of the simulation program with the same initialization parameter values to ensure that the values of the dependent variable are typical and not an outlier occurrence. Consider the simulation model introduced in chapter 2 that has nine parameters (see the details about the multiple restrictions on the possible discrete values for each simulation parameter in Table 2.4) with the total number of possible unique combinations of discrete parameter values 17,968,500. Furthermore, if we assume that 10 runs of the simulation will ensure a representative value of the dependent variable, then an analyst needs to create almost 180 million simulation runs to create an exhaustive parameter space. Assuming that each simulation lasts at least 2 seconds yields at approximately 4,159.375 days of computation. Most researchers would opt to reduce the number of discrete values or the number of repeated runs for each unique combination of values or split the workload among various workstations. While any of these solutions would reduce computational time dramatically, the former two would create less reliable conclusions and the latter one would require more effort on behalf of the researcher.

Sometimes, authors blame the overparametization for heavy computational requirements in simulation models (Fagiolo, Moneta and Windrum 2007). However, as Marks (2007) explains many simulation models are legitimately complex and require multi-parameter structures to accurately represent social life.[2]

The enormity of parameter spaces which make these models difficult to understand has been cited as one of the reasons for complex simulation models not becoming mainstream in social sciences (Fagiolo, Moneta and Windrum 2007; Axelrod 2006, 2003; Carley N.d.). In computational economics, the problem of analyzing simulation models has been separated into model verification and model validation. Verification refers to examining whether the software correctly implements the theoretical abstraction that the researcher intended (Galán et al. 2009). Validation refers to the degree that the model represents the phenomenon of social life that it was designed to capture (Galán et al. 2009). Validation could be measured in terms of goodness of fit to the characteristics of the empirical data Bianchi et al. (2007); Axelrod (2006); Carley (N.d.). The lack of common guidelines for describing, verifying, and validating simulation models has generated calls for establishing a standard for analyzing agent-based simulation models (Richiardi et al. 2006; Brenner and Werker 2007). It is important to point out that while verification needs to become a required step across social science disciplines, not all models are developed with a purpose to describe an aspect of social life accurately. Oftentimes, social scientists are interested in modeling hypotheticals to observed behaviors to emphasize what logically follows from one's assumptions (Primo and Clarke 2007). Thus, it would be preferable if the systematic analysis of computational models was not tied to how accurately they describe empirical reality.

[2] For a more formal treatment on the measure of complexity in simulation models, see Marks (2007, p.276-279).

On the other hand, the lack of strict guidelines in analyzing simulation models has also produced creative solutions of employing optimization heuristics to calibrate the model to match the empirical data (Bianchi et al. 2007), or conduct robustness checks to uncover which parameter values "break" the model's predictions (Miller 1998), or attempt to do both (Midgley, Marks and Kunchamwar 2007).[3] This chapter implements Miller's (1998) idea of employing optimization heuristics to generate and analyze a partial as opposed to exhaustive parameter space. Genetic algorithms can find global solutions of highly nonlinear relationships without generating an entire search space of the simulation model. This is possible because these heuristics rely heavily on the recombination of successful candidate solutions with a small degree of randomization to avoid being "stuck" in one region of the space.

I implement Miller's idea to infer relationships of interest among the model's parameters. The software program *SimGA* introduced in this chapter uses a genetic algorithm to search for the combinations of parameter values that minimize the distance between the simulation's output and the value of the "dependent" variable (a minimum or a maximum) specified by the analyst.[4] A partial parameter space is a subset of the exhaustive parameter space, which the algorithm has to search to locate the global solution. Since it takes less time to generate a subset of the parameter space than an entire search space, this method of interpreting simulation models is more efficient. As I decouple the simulation model from the genetic algorithm and from the code that handles information between the genetic algorithm

[3] Midgley, Marks and Kunchamwar (2007) use a genetic algorithm to "break" their model, but only discuss the possibility of using a genetic algorithm for calibrating an agent-based model to reproduce the empirical data generating process.

[4] I use the term "dependent" variable to refer to the emergent characteristic of the system. In the case of the benchmark model, the amount of armed conflicts or the amount of conflicts of interests or the average division of the disputed prizes at a certain point in time could all be viewed as emergent characteristics of the system.

and the simulation, this program can be used with any standalone simulation model that may be run in a batch mode and may generate numerical output.

3.2 Genetic Algorithms

Genetic algorithms (GAs) mimic the process of biological evolution by envisioning the candidate solutions as biological organisms represented by their genetic code (chromosomes). The fitness of an organism is analogous to the quality of a candidate solution. The program ensures that combinations of the "fittest" solutions (like breeding among the fittest organisms in nature) leads to desirable attributes in the next generation. In addition, genetic mutation ensures population diversity and thus provides a genetic algorithm with an important advantage over other optimization heuristics, as it avoids converging on local optima or becoming "stuck" in flat regions of the space.[5] Over time, due to selection, breeding, and mutation, organisms evolve to be increasingly fit. Thus, from one iteration to the next, GAs present better and better candidate solutions to the optimization problem until the global solution is reached (Givens and Hoeting 2005; Mitchell 1997).

While there exist many variants of genetic algorithms, all share certain key features. Every GA starts with some finite population P of candidate solutions a, which are randomly generated in the initial P_0.[6][7] Each solution (often referred to as "chromosome" to emphasize the biological origin of the algorithm's logic) is comprised of alleles.[8] Through an objective function, one may determine what value of the response variable any given a has.

[5]See Mitchell (1997, pp. 155-156) for a discussion of GA's advantages over hill climbers and other optimization heuristics when the search space is large, multimodal, and nonlinear.

[6]Most of the notation is adopted from Bäck, Fogel and Michalewicz (2000)

[7]Another possibility is to start with the analyst's "best guess" of what a solution may be, however, this may inhibit the analyst's understanding of how reliable the algorithm is (Reed N.d.)

[8]As I describe later in this paper, in my application each allele represents a parameter value.

In this application, a GA treats the user-specified minimum or maximum value of the dependent variable as the sought target T, which serves as the benchmark comparison for all solutions, such that each individual a is evaluated based on how closely it approximates T.[9] In other words, each candidate solution's fitness score θ equals the inverse of the distance between the sought target and the individual's value of the response variable, $\theta_a = \dfrac{1}{T - T_a}$. Each subsequent generation P^{λ} is selected from the highly scoring solutions of the previous generation P^{μ}, by pruning the worst performing half of P^{μ} (this selection method is often referred to as "elitist", Bäck, Fogel and Michalewicz (2000)). The most "fit" parent solutions are then used to produce the next generation through applying the crossover c and mutation m operators that mimic biological breeding and mutation in nature.[10] With probability ρ_c, the crossover operator cuts and swaps two parent strings at a random single point to generate two children strings (as shown in Table 3.1).[11] After crossover is complete, the size of generation P^{λ} is the same as the size of generation P^{μ} , as in addition to each pair of parent strings there also now exists a newly created pair of children strings. After selection and crossover, a mutation operator is applied to each of the strings in generation P^{λ}, as shown in Table 3.2. For a mutation operator, the algorithm loops through each allele of each individual and replaces its value with a randomly generated value with probability ρ_m. As Mitchell (1997, pp. 175-176) explains, the research on optimal parameter values for GAs has revealed that a population size P of 20-30 individuals, ρ_c of 0.7-0.9 and ρ_m of 0.005-0.01

---

[9] Here, I use a genetic algorithm with the parameter values based on the discussion of encoding a GA in Mitchell (1997, pp. 166-176).

[10] For the discussion on variations of crossover and mutation genetic operators, see Givens and Hoeting (2005); Mitchell (1997).

[11] It is important to note that some parameters in agent-based models may have narrower ranges of possible values than other parameters. Therefore, the crossover operator in *SimGA* does not change the index location of an allele, while other applications often apply such change.

per bit are optimal. Since I omitted double to bit conversion, I use a higher ρ_m of 0.3.[12] This process is repeated from one generation to the next. Since every next generation is selected based on the solutions' fitness scores, the individual distances between T and each individual's T_a become smaller with time, until the GA finds the solution that matches the target value.

To summarize, a basic GA sequence includes the following steps (see Figure 3.1 for more specific pseudocode and a graphical representation in Figure 3.2):

1. Make the first generation $P^{\mu}(0)$,
2. Check whether the fitness threshold is satisfied, i.e., $T_a \leq T$[13] or whether the maximum number of generations has been achieved, i.e., $t \geq t_{max}$[14]):
   - If yes, terminate.
   - If not, continue.
3. Prune population, select pairs to mate,
4. Generate population $P^{\lambda}(t+1)$ (using selected pairs),
   - Apply crossover operator with probability ρ_c
   - Apply mutation operator with probability ρ_m
5. Evaluate each a's fitness θ_a in $P^{\lambda}(t+1)$,
6. Go to step 2.

[12] As explained later in this paper, all parameter values in the example I use are encoded as double values with six decimal places, where each such double has a 64 bit representation. The probability of a double changing through the probabilistic mutation of each bit is $64 \times 0.005 = 0.32$. Omission of double to bit conversion has become commonplace in recent applications to preserve efficiency of the algorithm (Reed N.d.).

[13] If the sought global optimum is a maximum, then the condition is altered to be $T_a \leq (-1)T$.

[14] In the application, $t_{max} = 5,000$; this condition is necessary to terminate the program in the case the algorithm is "stuck" and the mutation operator is not able to move it away from the plateau region.

3.3 Analyzing Simulation Models with Genetic Algorithms

The software application described in this section provides communication between a genetic algorithm search program with simulation models, conceptualizing any simulation model as a fitness function of the genetic algorithm. In the context of analyzing an agent-based model, each candidate solution a is represented as a string of the model's initialization parameter values, while an objective function that links a chromosome to the response variable is the simulation model itself.[15] Therefore, for the GA to evaluate an individual solution's performance (i.e., to compute θ_a), we need to execute an agent-based model with that solution's initialization parameter values and record the numerical output from the simulation, which would provide us with the value of the "dependent" variable for a given set of the initialization parameter values. Due to the randomization of some steps, agent-based models are almost never deterministic. That is, to measure a value of the response variable for each individual chromosome, a researcher needs to execute a model n times and then calculate a mean/median of the response variable's outcome over n simulation runs. This ensures that a researcher evaluates a representative value of the response variable for each chromosome as opposed to an outlier value. Therefore, in the context of analyzing an agent-based model, θ_a equals the inverse distance between the target value of the response variable and the mean/median of the individual target values over n runs of the simulation, or $\theta_a = (T - \frac{\sum_{i=1}^{n} T_a}{n})^{-1}$.

[15] Miller (1998) initially proposed this approach and Midgley, Marks and Kunchamwar (2007) are the only other known to me implementations of this approach, although they do not make their code available for replication or user customization.

Since each candidate solution is a string of parameter values, the number of elements in each solution equals the number of varied parameters in the simulation as determined by the researcher. In the context of agent-based models, model parameters are the major characteristics of the model's agents or the system as a whole. GAs do not require any particular number of elements (alleles) in each chromosome, although it would be impossible to perform crossover for solutions with fewer than two elements. This flexibility in the size of individual solutions allows researchers to choose which parameters of the simulation model to vary and which ones to hold constant. In the following sections, I demonstrate the work of the algorithm with all varying nine element chromosomes, however, it is fairly simple to constrain certain parameters such that the algorithm is not allowed to evolve those elements of candidate solutions.

3.3.1 Pseudocode for the *SimGA* application

The proposed software application *SimGA* includes routines that (a) execute external simulation code, running the simulation model with the candidate solutions, or strings of initialization parameter values generated by the genetic algorithm, (b) score each string by analyzing the output from a batch of simulation model's runs, and (c) apply a simple genetic algorithm to the scored parameter strings. The code includes the following steps:

1. Set target value T of the dependent variable.
2. Randomly generate the initial population of strings of parameter values $P(0)$.
3. In population $P^{\mu}(t)$, for each string of parameter values a:
    - Write a parameter file,

- Run the external simulation model n times to generate n summaries of the numerical output,
- Compute T_a: calculate the mean value of the dependent variable across n runs,
- Compute θ_a: score a to reflect its fitness.

4. Record P candidate solutions to text file, such that each string of initialization parameter values has its corresponding fitness score and individual target value.

5. Check whether any of the scored strings' individual target values match the sought target (i.e., $T_a \leq T$) and whether the maximum number of generations have been created (i.e., $t \geq t_{max}$):[16]

   - If yes, the *SimGA* run terminates. If the run terminates and the SimGA application is set to run in batch mode, go to step 2.
   - If not, continue to step 6.

6. Apply the remaining GA steps:

   - Prune the worst fitting half of the population.
   - Make a new generation $P^{\lambda}(t+1)$:
     - Apply the crossover operator with probability ρ_c,
     - Apply the mutation operator with probability ρ_m.
   - Save $P^{\lambda}(t+1)$ as $P^{\mu}(t)$.
   - Go to step 3.

These routines are also shown graphically in Figure 3.3, which uses color coding to demonstrate the steps that involve the genetic algorithm (shown in blue) and the steps that involve interfacing a simulation model with the genetic algorithm (shown in green). In addition, the key features of the *SimGA* application are summarized in Table 3.3.

3.3.2 Illustrative run: the benchmark model as a fitness function

The *SimGA* application may be coupled with any external simulation model that generates numerical output, may be executed in a batch mode, and may be initialized with a

[16] See footnote 14 in step 2 of the genetic algorithm outline.

parameter file. Given that most simulation models would satisfy these criteria, a heuristics-based optimization of initialization parameter values is potentially a general solution to the large parameter space problem. This method may be used for both exploring major relationships among a model's parameters and for performing robustness checks. The application is written in the Java programming language.[17] Given the amount of computational power needed to run a simulation model thousands or hundreds of thousands of times, it is more efficient to conduct all the computations remotely via a high performing computing infrastructure. Therefore, the *SimGA* program is accompanied by three Bash scripts, the goals of which include: (i) to compile and initialize *SimGA*, (ii) to compile and initialize an external simulation model, and (iii) to request a remote batch execution of the *SimGA* application.

In this subsection, I discuss the sequence of steps within a single iteration of the *SimGA* software application, when it is coupled with the null agent-based model developed in chapter 2. The benchmark agent-based model represents an international system of dyadic intra- and post-war bargaining interactions. Recall that third parties do not affect actors' decision-making in the null model. The response variable of interest is the amount of armed conflicts experienced by the system as a whole. It is important to note that the goal of this exercise is to verify the ability of this method to evolve initialization parameter values in such a way that the model's output would match the sought fitness conditions. In chapter 2, I have demonstrated analytically and verified computationally that the benchmark model has two important takeaways: the greater the costs of fighting, the less likely a dyad is to fight and the greater the difference between A's probability of defeating two different types

[17] I also employ the RePast libraries to import a "canned" version of the parameter reader, however this library is not required, and may be substituted by developing a routine that parses through a text parameter file or it is also possible to import some other library to handle the task.

of B, the more likely any given dyad is to go to war. I will now explore whether *SimGA* is able to generate the same conclusions when communicating between the benchmark model and the genetic search algorithm.

Recall that the null model has two modes of generating numerical output. First, the model may save output data recorded as a time series of dyadic interactions. Such an approach allows one to model heterogeneous dyads and is preferable for the purposes of verifying computer code and understanding the details of intra-dyadic interactions. In contrast, when the output needs to be delivered as a single line summary of system level emergent properties, it is preferable to execute the model with "quasi-homogeneous" dyads, so that there would be no need to create a time series for each dyad. Given that the dyads are independent of each other by design, having all dyads in the system share the same features (e.g., distribution of capabilities or initial beliefs) is not a problem for inference, as it is possible to compare across systems. The second approach is preferable when the model is executed in batch mode, with the purpose of generating a representative parameter space to analyze the model dynamics. Although all pairs of actors share the same initial characteristics, the value of the dependent variable differs from one run to another due to randomization steps in the simulation. This is why the simulation model is executed n times with each set of initialization parameter values, i.e., candidate solution. Each run generates its own value of the response variable T_a. After n model runs, the average T_a is computed so that a given individual a's fitness score is representative across multiple runs, as opposed to being an outlier value. In this application of *SimGA*, the null model is executed in a batch mode with single line summaries of each individual run.

The application seeds the first generation of candidate solutions randomly.[18]Each generation has 50 strings of initialization parameter values (interchangeably referred to as "individuals" or "candidate solutions"). Tables 3.1 and 3.2 provide examples of candidate solutions for the benchmark model or, in the context of a GA, nine-allele chromosomes. All parameters are double values with six decimal places.[19] Given that each generation has a population of 50 candidate solutions and each candidate solution requires 50 runs of the simulation to generate a representative value of the dependent variable (here, the count of armed conflicts experienced by the system as a whole), each generation requires 2,500 function evaluations, i.e., simulation runs. To create each subsequent generation, the strings are ranked according to their fitness scores and the half worse performing strings are eliminated. The better performing half of the generation are then recombined and mutated as described in section 3.2.

Figure 3.4 shows the fitness scores for each generation over the course of one execution of *SimGA*. The best fitness score per generation is shown in red, while the average generation's fitness score is in blue. The initial seed is random. The termination conditions are either for the best individual T_a value to match the sought T or for t to surpass t_{max}. Even if one candidate solution achieves the desired fitness, *SimGA* starts a new execution. Recall that each fitness score is computed as the inverse distance between each candidate solution's average T_a value (over multiple simulation runs) and the sought T. In this example, the sought fitness constitutes 100 armed conflicts, as each pair of governments may

[18]A feature of seeding the "best guess" values as the initial generation may be easily implemented as well, however the issue of reliability, i.e., the ability of the algorithm to find global optima regardless of its starting position, may arise, as the algorithm may not be fully tested if the analyst employs the "best guess" values as initial populations.

[19]See chapter 2 for the details on the boundary conditions of each parameter and interconnectedness of some of the parameter values, specifically, Table 2.2.

only experience two conflict periods by design, therefore a system of 50 dyads may have the maximum of 50 conflict periods overall.[20] If a given string of initialization parameter values generates on average 99 armed battles, then fitness score is computed as 1. The graph does not show the fitness score for the last generation, because the fitness scores for the solutions that closely approximate the maximum exceed 1. This is by design: the fitness score formula is the inverse of the distance between the target sought and target achieved. Once the achieved target values become high enough the fitness score values become large, e.g., fit score for target value of 99.5 will be 1/(100-99.5)=2 and fit score for target value if 99.8 will be 5.[21] It is important to note that the null simulation model is non-deterministic, i.e., randomization steps within the simulation algorithm lead to different values of the response variable for each given execution of the simulation. Having on average 100 battles means -literally- that every single execution of the model yields 100 battles (given that the recording format allows for enough precision); in this application, I have allowed for the algorithm to round two digits after the decimal point, which often leads to averages of 99.95 and above to be assigned a T_a value of 100.

The x axis in Figure 3.4 shows the order of the generation; given that each generation has a population of 50 candidate solutions and each candidate solution requires 50 runs of the simulation to generate a representative value of the dependent variable, each generation requires 250 function evaluations, i.e., simulation runs. In the charted example, there were

[20] Recall that I assume that B may be weak, moderately strong, or strong, which means that B may reject two A's offers before A knows with certainty B's type. As battles are fought only when B rejects, a maximum of two battles may be fought in each dyad. See section 2.1 for more details.

[21] For those cases that score on average an exact value of the sought target, the *SimGA* algorithm prevents the division by zero and assigns a fitness score that is still a real number.

38 generations, therefore, the application executed $2500 \times 38 = 95,000$ simulation runs to achieve the sought fitness.

Figure 3.4 demonstrates the nonlinearity in the formula used to compute fitness scores: the closer an individual solution's T_a value to the sought target, the more it is weighted as a better solution. Figure 3.5 underscores this idea further by plotting the best (shown in red) and mean target values (shown in blue) for each generation over the course of the same single execution of *SimGA*. While the best target value achieved in generation 30 is 80 battles, which makes up 133.33% of the 20th generation's best T_a value of 60 battles, it is assigned a fitness score of 0.05, which constitutes 200% of the 20th generation's best fitness score of 0.025.

A few major takeaways follow from Figures 3.4-3.5. First, the algorithm improves both the average score and the best performing string from one generation to the next until the solution is found. Second, the rate of change in the best string is faster than the rate of the average improvement in each generation. The average result may deteriorate in $P(t+1)$ compared to $P(t)$ because the newly recombined and mutated "children" members may perform worse than their "parents," however, since I am preserving the best performing parent string unchanged, the algorithm is always able to get rid of the worse performing solutions by pruning them from the population. Third, the rate of improvement from one generation to the next is not constant, experiencing "jumps" (see best T_a values in generations 9-12) and periods of relative stagnation (see best T_a values in generations 26-30). Fourth, even though I am preserving the best performing half of the generation and evaluate an average value of the dependent variable across 50 simulation runs, it is possible for the best T_a value to decrease in the $P(t+1)$ generation. This is due to the nature of this exercise: agent-based models have multiple randomization steps which mean that the value of the dependent variable is not going to be exactly the same across multiple runs even though the initialization parameter values are the same. I have chosen to rely on the average value of the response variable across 50 runs as a compromise between attaining a representative value and the amount of time needed to compute it. Having *SimGA* execute the null model 100 times to score each candidate solution yields much less variation in the outcome of the dependent variable when the null model is seeded with the same initialization parameter values.

3.3.3 Results from *SimGA* compared to exhaustive search

While plotting individual runs' time series may provide a sense of how the algorithm performs over time, it is important to examine the relationships among the model parameters to ensure that the genetic search program evolves the initialization parameter values in such a manner that the relationships predetermined by the model's assumptions (section 2.1) are borne out in the *SimGA*'s output. Recall that each dyadic interaction in the benchmark model is independent of the rest of the international system. Two major conclusions come out of the null model's assumptions. First, the greater the costs of fighting, the less likely A is to risk war. Second, as the difference in the probability of A defeating two different types of B increases, A becomes more likely to risk an armed conflict.[22] In this section, I focus on the latter conclusion from the null model. I analyze the output from multiple runs of *SimGA* for the purpose of verifying if the same conclusion could be uncovered without conducting an exhaustive search of the null model's parameter space. The following discussion does not discover the said relationship between the difference in A's probability of defeating two

[22] See A's decision-making rules (expressed as inequalities 2.1-2.5 that detail these relationships.

different types of B and A's propensity to risk war, rather, it demonstrates that the search program finds the relationship predetermined by the model's assumptions.

The output from *SimGA* runs is recorded as numerical data, which include all generated candidate solutions in all populations over 500 runs of the *SimGA* application. Each observation in the data set consists of three major elements:

- a string of initialization parameter values for the null agent-based model,
- an average value of the dependent variable (here, the count of battles in the system) that this string generates across 50 executions of the simulation model,
- a fitness score computed based on the average of the dependent variable.

Thus, each observation is a summary of 50 simulation runs or function evaluations. Most executions of *SimGA* take between 20 and 60 generations to reach the desired fitness. The average number of generations is 56.7, the minimum is 6 and the maximum is 132. Figure 3.6 shows the kernel density plot of number of generations per single execution of *SimGA*, i.e., the number of generations it took the genetic algorithm to evolve the candidate solutions such that at least one solution in a generation reaches the sought fitness. The total number of observations in the data is 32,574, representing a total of 1,628,700 function evaluations. Since the termination conditions include reaching the target of 100 battles, the data include 500 strings that generate on average 100 battles (or rather, 99.995 battles or more that are rounded up to 100).

To analyze the data output from the genetic search application, an analyst may select to focus on interpreting the patterns within solutions only or to explore how the solutions differ from the rest of the population(s). Below, I present the latter version of analysis as the null model is simple and there is little variation within the solution strings. It is important to note, however, that more complex models would necessarily yield some variation among the solution strings, which implies that the software would benefit from an addition of an automated pattern recognition module as a future extension.[23]

To analyze the numerical output from *SimGA*, I have constructed the variables that measure the differences in A's probabilities of defeating two different types of B. The same variables were built in chapter 2, section 2.3.3. The measures are: $Difference_{\overline{p}_{t=1},\hat{p}_{t=1}} = \overline{p}_{t=1} - \hat{p}_{t=1}$. $Difference_{\hat{p}_{t=1},\underline{p}_{t=1}} = \hat{p}_{t=1} - \underline{p}_{t=1}$. $Difference_{\overline{p}_{t=1},\underline{p}_{t=1}} = \overline{p}_{t=1} - \underline{p}_{t=1}$. The simulation model implies that higher values of *Difference* increase A's probability of risking war and thus dyad's propensity to fight a war. Recall that I assumed that A's probability of winning a battle equals the proportion of the dyad's total capabilities that A possesses. If A and B's capabilities are similar, then the difference between A's probabilities of winning against two different types is greater than when A and B's capabilities are different. The more battles pairs of governments fight during a single simulation run, the closer the outcome variable approximates the sought target of 100 armed conflicts and the higher the fitness score of a given string. Just like in the illustrative run of section 3.3.2, each system in the batch mode has 50 dyads, each of which may fight a maximum of two battles, therefore the maximum possible number of battles fought in the system is 100.

Figure 3.7 presents three superimposed kernel density plots of different simulation outcome groups (high level of armed conflict in the system vs. low level). Blue color indicates those densities that correspond to high levels of armed conflict in the system (T_a values include 90-100 battles); red (0-10 battles) and green colors (10-20 battles) mark those

[23] See subsection 6.3.2 for more details.

densities that correspond to low levels of armed conflict in the system. The graphs demonstrate that higher levels of armed conflict occur when the difference between A's probability of defeating two different types of B is higher. While the densities of the poorly performing strings are indistinguishable, the densities of the (close to) solution strings do not overlap with the worse performing members of the population, which confirms that the search program has found exactly what it was expected to: the implied relationship between violent conflict and the difference in A's estimation of defeating two different types of opponents.

Table 3.4 presents Pearson correlation coefficients for differences in A's probability of defeating two different types of B and measures of candidate solutions' performance. The table shows highly significant correlation levels (in terms of both substance and statistical significance) between the differences in A's probability of defeating two different types of B and the candidate solution generating a high value of the simulation response variable (T_a) and thus being assigned a high fitness score (θ_a). Again, these high levels of correlation confirm that the *SimGA* program has been able to find the relationships implied in the null simulation model. In other words, this exercise verifies that the *SimGA* application is effective at achieving its goal of evolving the initialization parameter values of an agent-based model until candidate solutions meet a certain criterion of fitness.

Furthermore, *SimGA* required 1,628,700 function evaluations to generate enough numerical output for the researcher to establish the sought relationship. Recall that the same result was achieved by executing over 179,685,000 millions function evaluations when using an exhaustive parameter space approach outlined in section 2.3.3. This comparison underscores the efficiency with which genetic algorithms may deliver global solutions while only generating portions of the full parameter spaces.

3.4 Conclusion

This chapter introduces a software application designed to address a large parameter space problem by interfacing a genetic algorithm with a simulation model and treating a simulation model as a fitness function of the genetic algorithm. A large parameter space problem is at the root of the difficulties that analysts experience with verifying and validating simulation models. The proposed program relies on a genetic algorithm to evolve the initialization parameter values for simulation models until the performance of the model matches a certain criterion of fitness. The method is effective at achieving sought fitness criterion as it recombines the best solutions from the previous generations to improve performance and mutates some solutions' elements to avoid being stuck in a flat region of the space. The method is also efficient as it requires many fewer function evaluations to create a partial parameter space, still sufficient for inference.

This is not to say that the creation of a partial parameter space as opposed to an exhaustive parameter space requires no further interpretation. Statistical methods or automated pattern recognition algorithms could be employed to analyze the output from the genetic search software. In this chapter and in chapter 4, I employ statistical methods to understand the output from the genetic search program. While the null model is simple enough to assess with kernel density plots and correlation matrices, the extended model in chapter 4 demonstrates the need for an addition of more sophisticated interpretative tools to the genetic search program, which need to be incorporated in the future iterations of the program.[24] Given the advantages of the proposed program, the use of heuristics-based

[24] See subsection 6.3.2 for further details.

optimization techniques could become a more efficient substitute for an exhaustive search approach when analyzing simulation models.

Fig. 3.1. Pseudocode of the genetic algorithm

$t := 0;$
initialize $P^{\mu}(0) := a_1(0), ..., a_P(0)$
while$((T_a \leq T)||(t \geq t_{max}) \neq \text{true})$ **do**
  select:$P^{\lambda}(t+1) := s_{(\rho_s)}(P^{\mu}(t));$
  crossover: $P^{\lambda'}(t+1) := c_{\rho_c}(P^{\lambda}(t+1));$
  mutate:$P^{\lambda''}(t+1) := m_{\rho_m}(P^{\lambda'}(t));$
  score:$P^{\lambda'''}(t{+}1) :=$ for each a compute θ_a, $(P^{\lambda''}(t{+}1));$
$t := t+1;$
$P^{\mu}(t) := P^{\lambda'''}(t+1)$
od

Note: the pseudocode describes the specific genetic algorithm employed in this project as described in (Mitchell 1997, pp.166-176); the notation is based on Bäck, Fogel and Michalewicz (2000) and the outline is fashioned after the outline of a general evolutionary algorithm shown in (Coello Coello, Lamont and Veldhuizen 2007, p. 29). See Table 3.3 for reference.

Fig. 3.2. Flow chart of the genetic algorithm

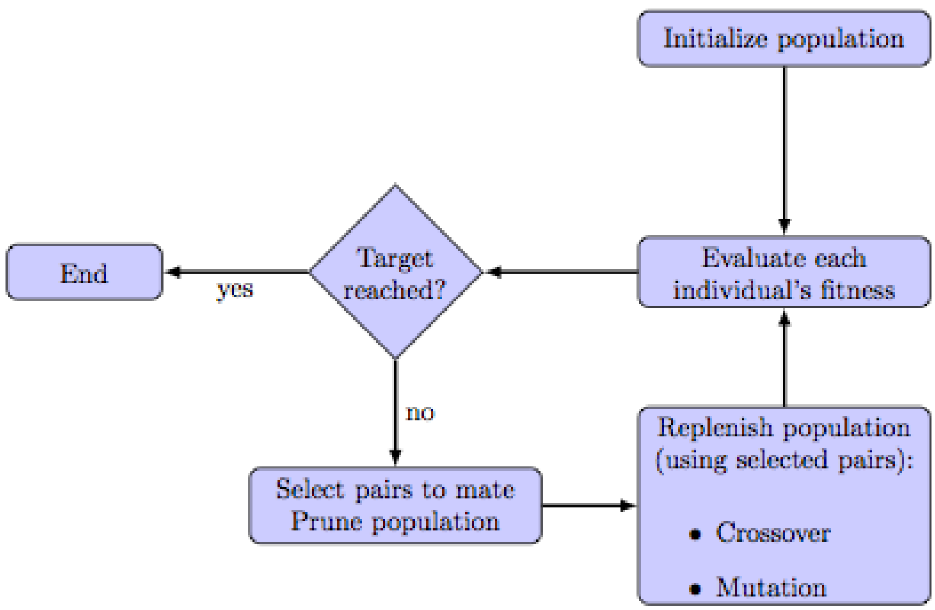

Note: the chart visualizes the logic of the genetic algorithm as described in Mitchell (1997, pp. 166-176).

Fig. 3.3. Flow chart of the automated search program, *SimGA*

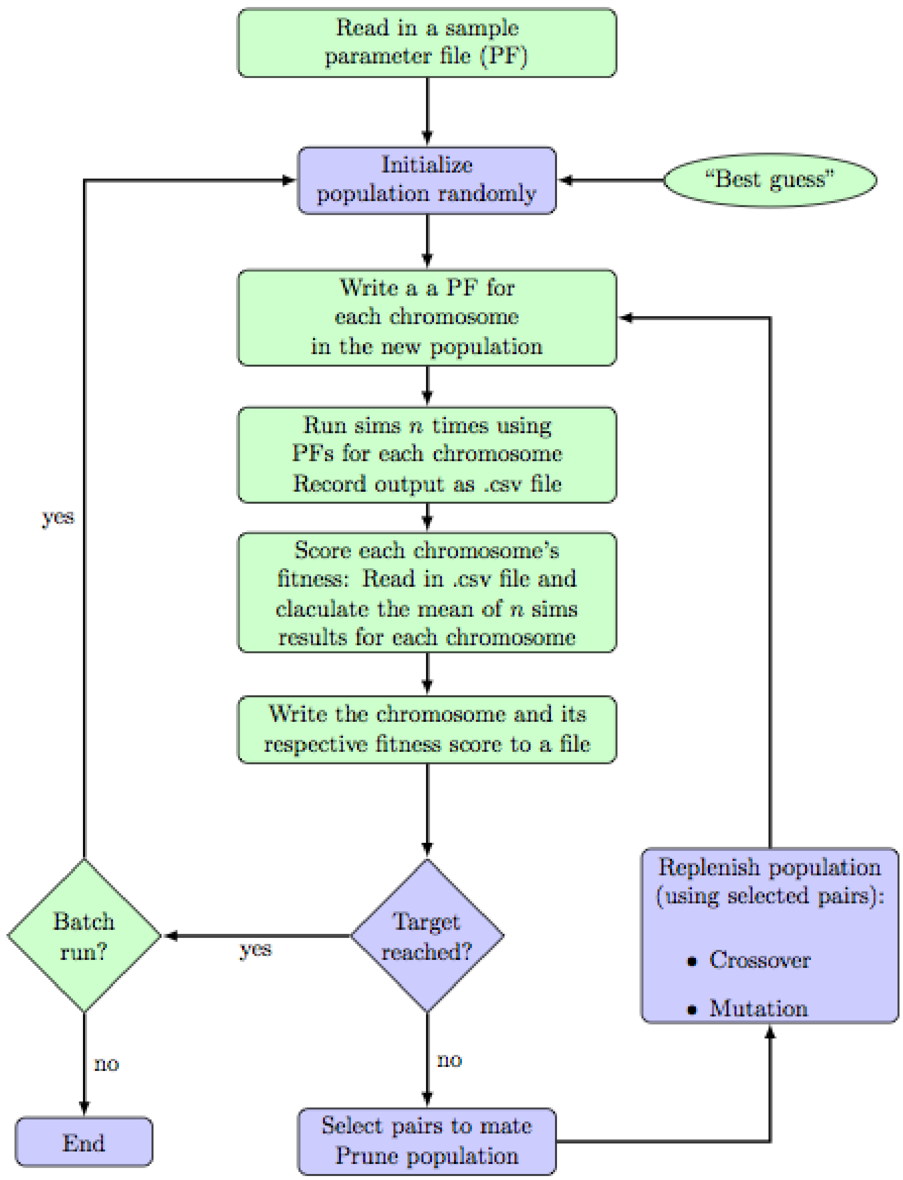

Note: the chart uses color coding to demonstrate the steps that involve the genetic algorithm (shown in blue) and the steps that involve interfacing a simulation model with the genetic algorithm (shown in green).

Fig. 3.4. Time series of candidate solutions' fitness scores over the course of a single complete execution of *SimGA*

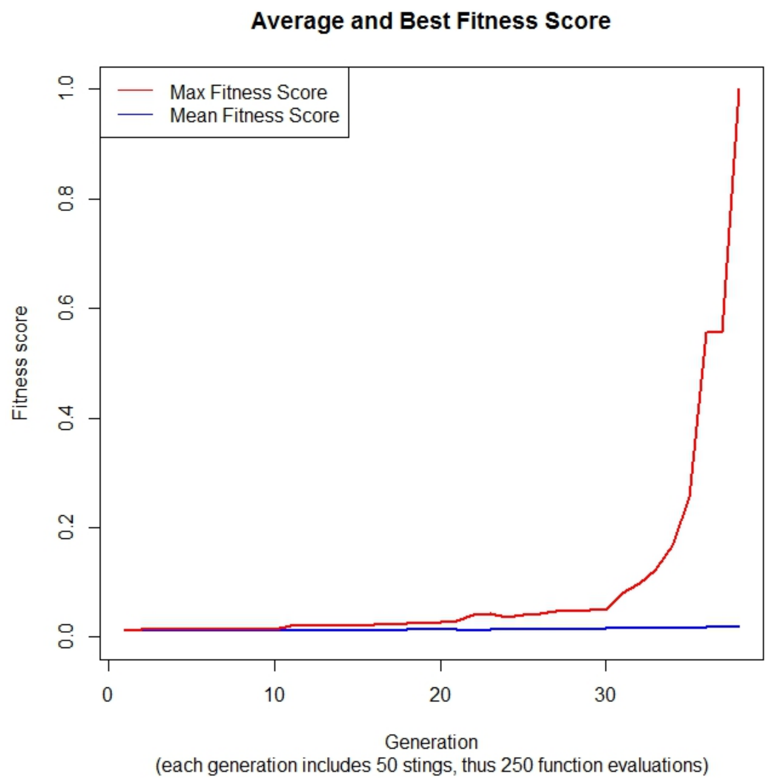

Note: best fitness score per generation is shown in red, average generation's fitness score is shown in blue. The initial seed is random. The termination conditions are for best individual T_a to match the sought T or for t to surpass t_{max}. Even if one candidate solution achieves the desired fitness, *SimGA* starts a new execution. The x axis shows the order of the generation; given that each generation has a population of 50 candidate solutions and each candidate solution requires 50 runs of the simulation to generate a representative value of the dependent variable, each generation requires 2,500 function evaluations, i.e., simulation runs. In this example, there were 38 generations, therefore $2500 \times 38 = 95,000$ simulation runs to achieve the sought fitness. The y axis shows the fitness score value, computed as $\theta_a = (T - \frac{\sum_{i=1}^{50} T_a}{50})^{-1}$. The graph does not show the fitness score for the last generation, because the fitness scores for the solutions that closely approximate the maximum exceed 1. This is by design: the fitness score formula is the inverse of the distance between the target sought and target achieved. Once the achieved target values become high enough the fitness score values become large, e.g., fit score for target value of 99.5 will be 1/(100-99.5)=2 and fit score for target value if 99.8 will be 5.

Fig. 3.5. Time series of generations of candidate solutions' T_a values in a single complete execution of *SimGA*

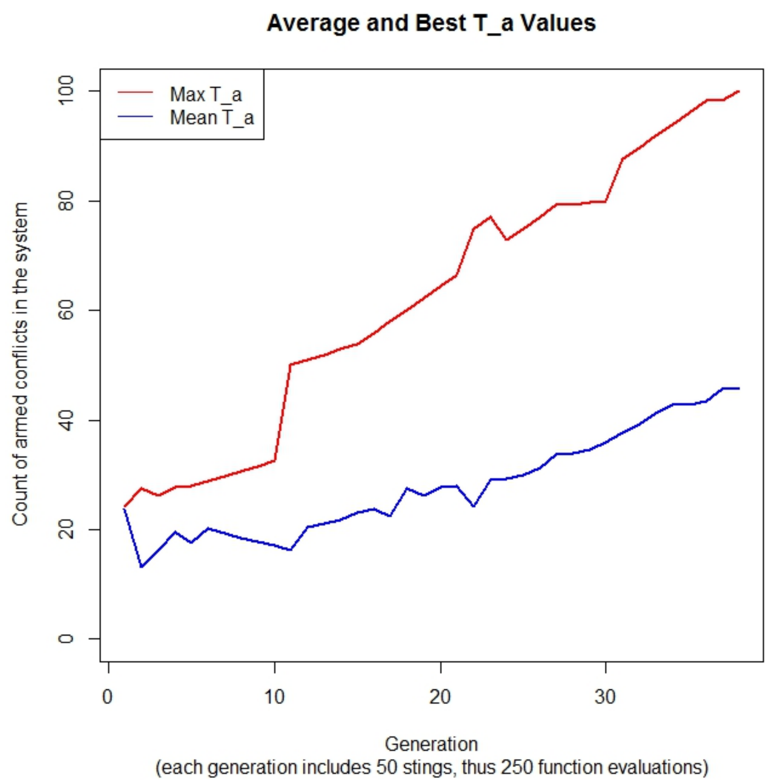

Note: best T_a per generation is shown in red, average generation's T_a is shown in blue. The initial seed is random. The termination conditions are for best individual T_a to match the sought T or for t to surpass t_{max}. Even if one candidate solution achieves the desired fitness, *SimGA* starts a new execution. The y axis shows the average value of the dependent variable (here, the count of armed conflict in the system) from the simulation numerical output. The x axis shows the order of the generation; given that each generation has a population of 50 candidate solutions and each candidate solution requires 50 runs of the simulation to generate a representative value of the dependent variable, each generation requires 250 function evaluations, i.e., simulation runs. In this example, there were 38 generations, therefore $2500 \times 38 = 95,000$ simulation runs to achieve the sought fitness.

Fig. 3.6. Kernel density of number of generations per single execution of *SimGA*

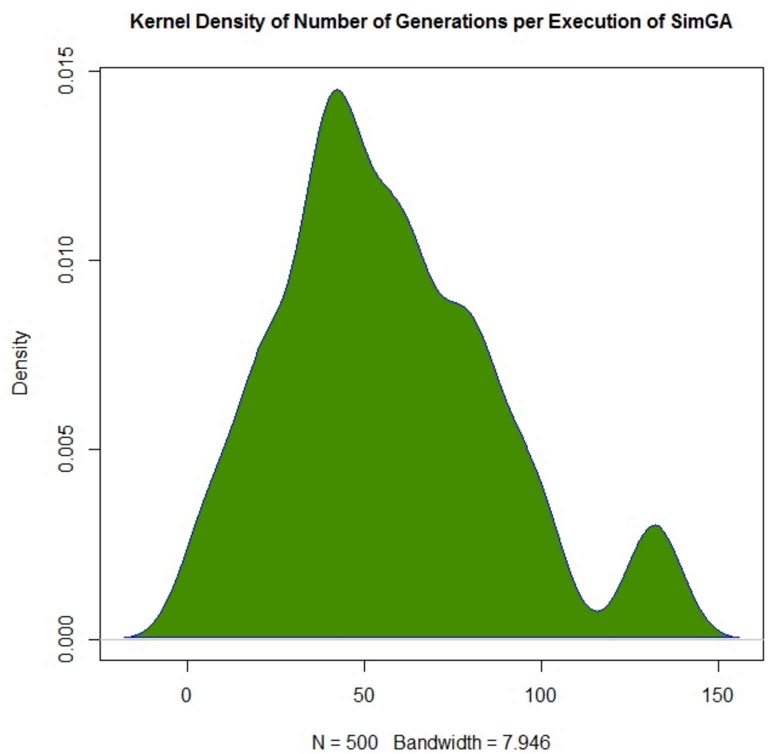

Note: the plot shows the kernel density distribution of the number of generations per one execution of *SimGA* program, i.e., the number of generations it took the genetic algorithm to evolve the candidate solutions such that at least one solution in a generation reaches the sought fitness. The average number of generations is 56.7, the minimum is 6 and the maximum is 132.

Fig. 3.7. Superimposed kernel density plots of different simulation outcome groups (high level of armed conflict in the system vs. low level)

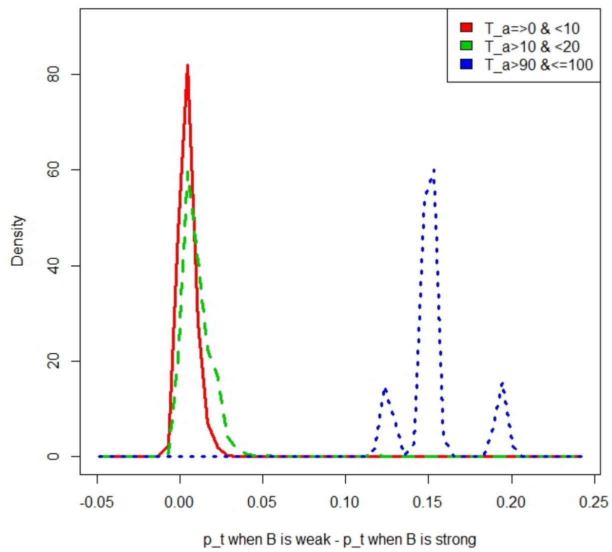

(a) x-axis is the difference in A's probability of defeating weak B and strong B

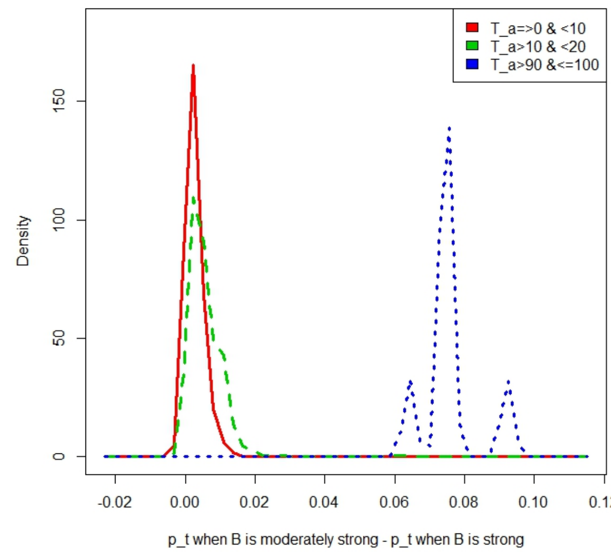

(b) x-axis is the difference in A's probability of defeating moderately strong B and strong B

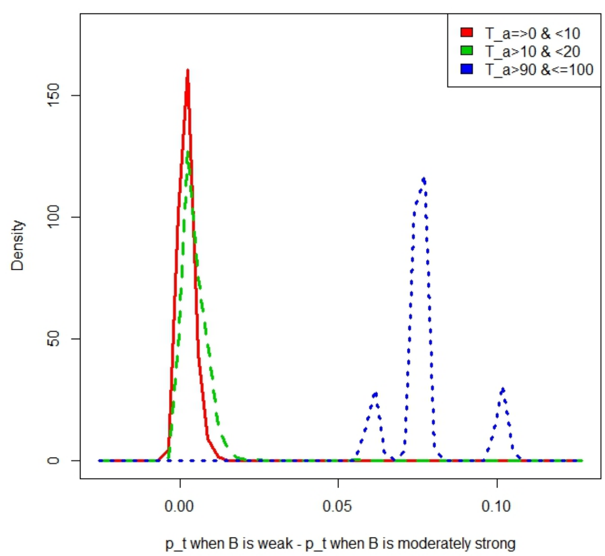

(c) x-axis is the difference in A's probability of defeating weak B and moderately strong B

Note: blue color indicates those densities that correspond to high levels of armed conflict in the system; red and green colors mark those densities that correspond to low levels of armed conflict in the system. The graphs demonstrate that higher levels of armed conflict occur when the difference between A's probability of defeating two different types of B is higher.

Table 3.1. Crossover genetic operator

Param Name	Param Value			
	1	2	3	4
	"Parent" Strings in $P^{\mu}(t)$		"Child" Strings in $P^{\lambda}(t+1)$	
m_A at t=1	*0.242797*	**0.985671**	*0.242797*	**0.985671**
m_B at t=1	*0.150054*	**0.893421**	*0.150054*	**0.893421**
ϕ_w at t=1	*0.082402*	**0.139076**	*0.082402*	**0.139076**
ϕ_m at t=1	*0.896749*	**0.789243**	**0.789243**	*0.896749*
$\underline{b}$	*0.016573*	**0.003568**	**0.003568**	*0.016573*
$\hat{b}$	*0.067785*	**0.015989**	**0.015989**	*0.067785*
$\overline{b}$	*0.074379*	**0.039742**	**0.039742**	*0.074379*
c_A	*0.010078*	**0.010078**	**0.010078**	*0.010078*
c_B	*0.025621*	**0.025621**	**0.025621**	*0.025621*

Note: With probability ρ_c, the crossover operator cuts and swaps two "parent" strings at a random single point to generate two "children" strings. It is important to note that some parameters in agent-based models may have narrower ranges of possible values than other parameters. Therefore, the crossover operator in *SimGA* does not change the index location of an allele, while other applications often apply such change. In addition, some of the parameters are labeled as a certain value at t=1, while other parameters do not have time specifications. This is an artifact of the null model: such characteristics as concealed capabilities $(\underline{b}, \hat{b}, \overline{b})$ and costs of fighting (c_A, c_B) do not change over time. In contrast, beliefs and capabilities of agents evolve as bargaining takes place.

Table 3.2. Mutation genetic operator

Param Name	Param Value	
	1	2
	Original String in p_t	"Mutated" String in p_{t+1}
m_A at t=1	*0.242797*	*0.242797*
m_B at t=1	*0.150054*	***0.271932***
ϕ_w at t=1	*0.082402*	*0.082402*
ϕ_m at t=1	*0.896749*	*0.896749*
$\underline{b}$	*0.016573*	*0.016573*
$\hat{b}$	*0.067785*	*0.067785*
$\overline{b}$	*0.074379*	***0.096459***
c_A	*0.010078*	*0.010078*
c_B	*0.025621*	*0.025621*

Note: For a mutation operator, the algorithm loops through each allele of each individual candidate solution and replaces its value with a randomly generated value with probability ρ_m. Also, some of the parameters are labeled as a certain value at t=1, while other parameters do not have time specifications. This is an artifact of the null model: such characteristics as concealed capabilities ($\underline{b}, \hat{b}, \overline{b}$) and costs of fighting (c_A, c_B) do not change over time. In contrast, beliefs and capabilities of agents evolve as bargaining takes place.

Table 3.3. Features of *SimGA* software program and the genetic algorithm embedded in it

Notation	Code reference	Interpretation	Default Value‡ (if applicable)
Parameters in *SimGA* set by the user, constant across generations			
T	`target`	The target value of the dependent variable	100
n	`numRuns`	Number of simulation runs necessary to to evaluate the fitness of one string a in generation t	50
P	`numStrings`	Population, i.e. the size of each generation	50
ρ_c	`prob_cross`	Probability, with which a crossover genetic operator is applied to a pair of strings	0.8
ρ_m	`prob_mutatn`	Probability, with which a mutation genetic operator is applied to a string	0.3
Quantities that are calculated for each solution in each generation			
T_a	`target_i`	Numerical output from the simulation; The value of the "dependent" variable in a single simulation run	NA
θ_a	`fitScore`	Fitness score for individual candidate solution a; $\theta_a = (T - \frac{\sum_{i=1}^{n} T_a}{n})^{-1}$	NA
Other notation			
t	`tick`	Order of the generation	NA
t_{max}	`tick max`	Maximum allowed order of the generation	5,000
a	`string`	An individual candidate solution within P	NA
μ	`parent`	A subscript that denotes "parent" string	NA
λ	`child`	A subscript that denotes "child" string	NA
c	`crossover`	Crossover operator	NA
m	`mutate`	Mutation operator	NA
N	NA	The count of all simulation runs necessary to generate the partial parameter space: the number of simulation runs for each member of the population in each generation, i.e., $N = n \times P \times t$	NA

Note: NA = not applicable. ‡ Default values are shown only for those parameters that are the same for each candidate solution. The indicated values were used in the example run discussed in section 3.3.2

Table 3.4. Pearson correlation coefficients for differences in A's probability of defeating two different types of B and measures of candidate solutions' performance

	θ_a	p-value	T_a	p-value
$Difference_{\overline{p},\hat{p}}$ at $t = 1$	0.81	0.000	0.78	0.000
$Difference_{\hat{p},\underline{p}}$ at $t = 1$	0.72	0.000	0.71	0.000
$Difference_{\overline{p},\underline{p}}$ at $t = 1$	0.76	0.000	0.74	0.000

Note: N=32,574. The table shows highly significant in terms of substance and statistical significance correlation levels between the differences in A's probability of defeating two different types of B and the candidate solution generating a high value of the simulation response variable (T_a) and thus being assigned a high fitness score (θ_a)

Chapter 4

An Extended Agent-Based Model of Conflict Management by Third Parties

In chapter 2, I develop and analyze the benchmark model of conflictual and peaceful interactions in the world without third parties. Since no third parties alter the distribution of capabilities or the costs of fighting in the middle of an ongoing conflict, the belligerents settle once no uncertainty remains about their relative military power. Therefore, the null model establishes that -when third parties do not interfere- war is an effective instrument for solving conflicts of interests and redistributing the disputed goods in accordance with the belligerents' relative capabilities. The benchmark model represents a stylized version of the world in which Edward Luttwak's (1999) plea to "give war a chance" was answered. Peace deals never fail in the null model, because no sudden changes in costs of fighting or external shocks to capabilities take place.[1] The extended version of the model, which I present in this chapter, portrays a more realistic snapshot of international relations: the sequences of conflict and peace that may be influenced by external actors through external support or power mediation. Conflicts recur regularly in the extended model and the condition of global peace is never achieved. I focus on how third parties worsen or alleviate the information problem in the dyads through conflict management techniques. s

In the extended model, I first explain why I group third-party behaviors as pressure for peace (also referred to as "power/coercive/leverage" mediation, which imply mediation

[1] An important extension to the null model would be to introduce random shocks in players' capabilities that simulate the discovery of natural resources or advances in technology.

that exerts pressure on the belligerents) and provision of external support. Combining these two groups of third-party behavior in the same model allows me to explore the effects of third-party influence on conflict processes in greater detail than previously done. I review the current state of research on the effects of third-party influence on conflict and model the agents' decisions to supply such influence relying on previous scholarship. I describe the extended model's additional actors (compared to the null model's actors) and their characteristics, pseudocode, and an illustrative run of the extended model in the subsequent section. I also focus on comparing the analysis of the extended model's parameter space by employing the genetic automated search to that of the null model.

4.1 Theoretical Foundations of the Extended Model

4.1.1 Two categories of third-party behaviors

Given the variety of third-party behaviors, a modeler faces a tradeoff between describing a representative range of actions and keeping his or her model tractable. I group third-party actions by their observable immediate intent while ignoring the long-term goals that third-parties may be pursuing. The focus on the observable actions allows me to group third-party behaviors as either attempts to keep both sides from fighting or attempts to reinforce one of the sides in the dyad. True motivations of third-party actions are difficult to infer from empirical reality. In contrast, the immediate intention is rarely disputable. For instance, pressure for peace normally seeks to stop violence in its immediate intention, yet may or may not have a long-term goal of favoring one of the sides over another. In contrast, providing weapons or troops to one of the sides intends to reinforce the recipient and thus achieve a more favorable settlement at the end of the war than what the recipient would have achieved without such support. It is unclear whether such support may also pursue a long-term goal of stable peace. As a result, I only focus on the observable actions: attempts to end violence or attempts to reinforce one of the belligerents.

Empirical examples of third-party actions that may fall into either of these categories is abound in the published data sets. Werner and Yuen (2005)[2] and Beardsley et al. (2006)[3] code pressure for peace in the interstate wars and interstate crises respectively. Furthermore, Beardsley et al. (2006) differentiate between facilitation of negotiations and manipulative mediation. In my model, I represent power mediation that seeks to pressure the belligerents into settling (prematurely) as a third party inflating the costs of fighting for the belligerents. Regan (2002)[4] codes provision of troops, naval forces, equipment or military aid, intelligence or advisors, and air support as various manifestations of military intervention. In addition the author measures provision of a grant, loan, credit, relief from past obligations, nonmilitary equipment as economic intervention. All said actions are represented in my model as transfers of resources from a third party to a given belligerent and I refer to these behaviors as reinforcements. f Although mediation and intervention literatures have evolved in isolation from each other, my modeling of external participation in conflict as either support- or mediation-providing actions is not unprecedented. Regan and Aydin (2006) have merged these two literatures by accounting for both reinforcement and diplomatic actions in the same model. Regan and Aydin (2006) label third-party actions as either "attempt[s] to influence the structure of the relationship among combatants," or "attempt[s] to manipulate

[2]The time domain of the pressure for peace variable was extended by Lo, Hashimoto and Reiter (2008).
[3]The International Crisis Behavior data were also used in Beardsley (2008, 2011).
[4]Further expanded by Regan, Frank and Aydin (2008).

the information that these actors hold." Regan and Aydin's first group is analogous to my description of support-providing actions. Reinforcing one of the sides in the warring dyad shifts the distribution of power in the dyad or in the authors' terms, alters "the power structure" of the conflictual dyadic relationship.

Regan and Aydin's second "diplomatic" group may appear analogous to my description of mediation-providing actions, however the mechanism assumed in Regan and Aydin differs from the mechanism I model. The authors' description of mediation as attempts to manipulate information that the belligerents hold hinges on the assumption that mediators have information that is unknown to at least one of the belligerents. Some of the current research challenges this assumption by pointing out that there has been no clear explanation of why a mediator would have information that is private knowledge and thus it is problematic to assume that mediators could improve upon unmediated negotiations (Kydd 2003; Fey and Ramsay 2010).[5] Due to our limits in understanding of how exactly information-manipulating mediation works, I choose to focus on mediation-like actions that operate through inflating the costs of fighting.

Instead of removing uncertainty, mediators may choose to threaten punishment or use leverage to pressure the belligerents into ceasefire. By using negative inducements (e.g., threats of intervention or threats of sanctions) and/or positive inducements (e.g., provisions of aid in exchange for peace), outsiders may inflate the costs of fighting so that violence appears a prohibitively expensive choice to the belligerents. (Bercovitch and Gartner 2006,

[5] I should note, however, that other research (Meirowitz et al. N.d.) derives that information-altering mediation can improve upon unmediated negotiations by, first, receiving private messages from the belligerents and, second, providing this information to the belligerents in the format of procedural mediation, i.e., mediation in which the rules of negotiations are set up by a third party, which limits the belligerents' negotiation options. Horner, Morelli and Squintani (N.d.) derive a similarly optimistic view of mediation that manipulates information. Still, neither Meirowitz et al. (N.d.) nor Horner, Morelli and Squintani (N.d.) specify which third parties would be trusted with private information and, more importantly, which mediators could be trusted by both sides of the conflicts.

p. 339) describe this mediation strategy as "[...]the most powerful form of intervention.[6] Here a mediator affects the content and substance of the bargaining process by providing incentives for the parties to negotiate or by issuing ultimatums. Directive strategies aim to change the way issues are framed and the behavior associated with them." Such forceful mediation is often labeled as "manipulative," "directive," or "power" mediation.

In sum, relying on the observable immediate intent of third-parties, I group external participation in conflict as either attempts to reinforce one side of the conflict or as attempts to end violence (the theoretical precedent was set by Regan and Aydin 2006). There is little disagreement about the mechanism through which reinforcement affects violence: the recipient of net surplus of support experiences a relative improvement of his/her position on the battlefield. In contrast, the exact theoretical mechanism through which third parties could change the asymmetric information between the belligerents is unclear. Thus, unlike Regan and Aydin (2006), I focus on the better understood type of mediation, which does not attempt to convey new information or establish new rules of communication, but instead makes fighting prohibitively expensive for the belligerents through exerting leverage and is referred to as power mediation.

4.1.2 Extant literature on power mediation and reinforcement

The effects of power mediation on war and peace duration: Conflict management research has recently established an important link between the strategy of third-party mediation and conflict duration. Theoretically, researchers had been distinguishing between

[6] It is not unusual for researchers to refer to mediation as a form of "intervention." Although the term intervention is more commonly used to refer to support-providing actions, some authors use it to describe any type of outside influence on conflict processes.

facilitative and directive mediation styles before large-N data became available by pointing out that a coercive mediation style exerts more leverage and changes the bargaining dynamic between the belligerents (Fisher 1995). Wilkenfeld et al. (2003) published the first test of whether mediation style affects crisis outcome and distinguished among three mediation styles: facilitation (i.e., channeling of information, "consultation model"), procedural (i.e., mediator has control over the rules of communication and publicity of the meeting), and manipulation (i.e., exerting leverage on the participants and changing their cost-benefit calculus).[7] The authors find support for the claim that manipulative mediation is the most successful mediation style at delivering the end of violence compared to more restrictive mediation styles. This result has been replicated across multiple data sets on international crises (Beardsley et al. 2006; Bercovitch and Gartner 2006; Beardsley 2011) and civil conflicts (Quinn et al. 2013).

Power mediation takes place in about 10% of the total number of all crises (Beardsley et al. 2006) and in 13.83% of the interstate wars (Lo, Hashimoto and Reiter 2008),[8] while making up 29-39% of the mediated cases (Beardsley et al. 2006; Bercovitch and Gartner 2006).[9] Substantively, power mediation increases the probability of ceasefire and a formal agreement by 28% when compared to the use of procedural mediation and by 60% when compared to unmediated crises among recognized governments (Wilkenfeld et al. 2003).[10]

[7] It is important to note that some studies omit the "procedural" category: "mediators may have access to more information than either disputant or possess sufficient resources to change disputants' preferences" (Terris and Maoz 2005, p. 565).

[8] It is important to note that in the subsample that encompasses the Cold War years, power mediators imposed ceasefires in 29.17% of cases (Werner and Yuen 2005).

[9] In the International Crisis Behavior data, 39.09% of mediated cases were coded as manipulation (Beardsley et al. 2006), while in the Bercovitch 2000 International Conflict Management data set, 29.10% of mediated cases were coded as directive mediation (Bercovitch and Gartner 2006).

[10] Wilkenfeld et al. (2003) find that facilitative mediation has no effect on the probability of signing a formal agreement.

In addition, Bercovitch and Gartner (2006) demonstrate that the higher the casualty rate in interstate crises, the more likely manipulation is to succeed at ending violence. In civil wars, manipulation leads to a 10-24% increase in the probability of signing a formal agreement compared to the facilitative mediation style (Quinn et al. 2013). It is clear that the success of power mediation is the result of a different mechanism of influence on belligerents that manipulators employ. Empirical papers construct mental models that do not require the same level of precision in describing a given mechanism as formal models do. This is why game-theoretic papers have struggled to explain how exactly facilitative mediation operates (Kydd 2003; Fey and Ramsay 2010), although there has been success in laying out formally the logic behind the mechanism of procedural mediation (Meirowitz et al. N.d.). Given the clarity of how manipulation operates, I model power mediation efforts as third-party attempts to make fighting less attractive for the belligerents through increasing their costs of fighting and do not take into account procedural or facilitative mediation.

Manipulation delivers peace – this is the takeaway that could seem obvious to policymakers given the aforementioned research. Yet, the second important development in conflict management literature in the past decade appears to have established an important caveat to power mediation's success at peace making. In continuation of bargaining tradition, researchers have argued that the style of mediation affects the outcome of conflict that, in turn, shapes the stability of the peace that follows. Werner and Yuen (2005) and Beardsley (2008, 2011) find that power mediation is more likely to generate unstable ceasefires after international wars and crises. In the literatureregarding post-civil war peace duration, Quinn et al. (2013) find that power mediation has no statistically significant effect on short-term post-conflict reduction of tensions.[11] Furthermore, DeRouen, Bercovitch and Wallensteen (N.d.) find that directive manipulation has a positive, albeit not statistically significant, effect on civil war recurrence. Although the findings in the post-civil-war peace literature are not conclusive, they do not contradict the findings of the post-interstate-war peace duration research. In other words, the empirical studies suggest that power mediators may choose peace in the short term over long-term stability.

The idea that powerful third parties, who have enough leverage to force the belligerents into ceasing fire, may create unstable ceasefires, is an important finding. Luttwak (1999) argues that third parties are not interested in building long-term peace because they are only motivated to stop "violence on [the] TV screens" of their domestic audiences. Walter (1997) argues that third-party security guarantees are a critical factor in achieving civil war settlement. Although Walter's analysis does not examine peace duration after ceasefire, a logical continuation of this argument would be that continued security guarantees also stabilize peace once the sides have signed a ceasefire agreement. Fortna's (2004, 2008) work on the effect of peacekeeper presence on ceasefire duration is an example of examining how third parties may continually create barriers to fighting even after peace is established.

The supply of power mediation: Why do third parties choose to mediate? Terris and Maoz (2005) and Melin (2011) describe a generic cost-benefit function for a potential mediator weighing her options: the expected utility increases as third party's benefit from a peace agreement rises, weighted by the likelihood that such an agreement is signed, and decreases as the cost of mediation becomes greater. Both studies focus on understanding what

[11] Although the effect is negative, suggesting that power mediation reduces tensions in the short-term after imposing a ceasefire on the participants of civil conflict.

determines mediators' preferences. Melin (2011) demonstrates empirically that the greater utility from agreement (measured as alliances with the belligerents, geographic proximity, colonial ties, and level of democracy within the warring dyad, given a third party is also a democracy) and the decline in cost of mediation (measured as greater capabilities of a potential mediator) make attempts at management more likely in interstate crises.[12] Terris and Maoz (2005) find that increased utility (measured as alliance ties and level of democracy, given a potential mediator is also a democracy) and a third party's ability to make a difference (measured as the preponderance of power within a dyad, i.e., the further the sides are from parity the more likely the stronger side to win) increase the incidence of intrusive mediation.[13]

In the civil war research, mediation is also more likely to be offered when colonial history, alliance ties, and geographic proximity are present; all these features may be viewed as measuring increased utility from forcing the belligerents to achieve a ceasefire (Greig and Regan 2008). Furthermore, Greig and Regan (2008) find that a third party having trade interests in a country experiencing a civil war is associated with an increased probability of that third party offering to mediate a settlement among the civil war participants. It is important to note that Greig and Regan (2008) do not distinguish between the strategies of mediation, nor do they control for third party capability.

[12] While Melin (2011) investigates attempts at conflict management, she describes these attempts as increasing economic pressure and exerting diplomatic or verbal leverage on the belligerents, therefore it is reasonable to compare Melin's dependent variables to what other researchers describe as power mediation.

[13] It is important to note that the major takeaway from Terris and Maoz (2005) is that the greater the versatility of the conflict (i.e., the diversity of games into which a given conflict game could be transformed), makes intrusive mediation more likely. Given that my model cannot be reduced to a simultaneous one-shot simple game, I abstain from making inferences from this aspect of the work.

I rely on these insights and model the general expected utility function after Terris and Maoz (2005) and Melin (2011) by letting the expected utility of the third party increase with her benefit from the end of violence and decrease with her costs of mediation.

The effects of reinforcement on war duration: The questions of how the type of outside support and the recipient of reinforcement affect civil conflict duration have received a lot of attention in the civil war "subsection" of conflict research. Patrick Regan contributed greatly to this line of work by introducing a detailed data set that records military support as provision of troops, naval forces, equipment or military aid, intelligence or advisors, and air support, and measures economic support as loans, non-military equipment or expertise, credits, relief of past obligations in post-WWII civil conflict (2000, 2002). Regan's own research does not take full advantage of these rich data by aggregating military and economic forms of external support into dichotomized indicators. Nonetheless, these data provide far more detail on the frequency and type of third-party activity than any other data set on external support in intra- or interstate wars.

Out of 150 recorded conflicts in Regan's (2002) data, there are 92 cases that received some sort of military reinforcement on either side. More specifically, there are 66 wars that experienced at least one instance of military support on the side of the government and 59 wars in which rebels received outside reinforcement at least once. Most conflicts that underwent outside reinforcements had many more than one instance of external support: only 10 wars had one instance of government support and 15 wars received support for rebels only once. The data record instances of outside actors sending troops (273 cases in support of government side and 179 cases to bolster rebels), providing weapons and equipment (82 cases of helping government advance and 74 cases to propel rebels), providing military training and expertise (63 cases) to improve the level of preparation of government troops and 19 instances of sending navy (13 observations to reinforce government and 9 cases to assist rebels), providing aircraft (104 cases to back government and 61 cases to further rebels).

Importantly, the author has collected information about conflicts with 200 fatalities and higher, including 27 cases in his data that do not fit the Correlates of War Project's definition of war that requires 1,000 battle deaths. Comparing the patterns of external support in Regan's data to those in the interstate militarized disputes data, one could readily compare the incidence of third parties providing troops to bolster either side in the conflict, or joining. In violent international militarized disputes, there were 395 instances of joining in the total of 1,660 militarized disputes (i.e., 23.79% of cases experienced joining) (Joyce, Ghosn and Bayer 2013).[14] Among civil conflicts, 79 out of 150 civil wars (i.e., 52.67%) had at least one instance of a third party sending troops (Regan 2002). While the difference in frequencies appears dramatic, it is important to note that the two data sets include various proportions of cases that did not escalate to a full definition of war. When comparing the incidence of joining in civil and interstate wars that generated at least 1,000 battle fatalities, the frequencies across war "types" appear similar: joining took place in 68 wars out of 123 civil wars (i.e., 55.28%) vs. 54 out of 107 interstate wars (i.e., 50.46%).

Power mediation occurs in far fewer cases than joining. Recall that the incidence of power mediation in international wars and civil conflicts is 10-14%,[15] making joining a

[14] By "violent" the authors refer to the disputes with 4th and 5th levels of hostility. Regardless of the hostility level, the MID data contain 2,332 disputes (Ghosn, Palmer and Bremer 2004).

[15] The sample in (Quinn et al. 2013, p. 23) contains low and high intensity intrastate conflicts in Africa in 1990-2005, as identified by the Uppsala Conflict Data Program (UCDP).

common event experienced by at least every second high intensity war regardless of whether the actors are recognized or unrecognized governments.

In terms of the effects of outside support on conflict duration, there have been inconsistent findings in the civil war literature. Regan and colleagues conclude that the aggregated indicators of external support extend the duration of civil wars (Regan 2000, 2002; Regan and Aydin 2006). Two studies – Balch-Lindsay, Enterline and Joyce (2008)[16] and Collier, Hoeffler and Soderbom (2004)[17] – have found that dichotomized measures of external support shortens civil war duration. However, the finding in Balch-Lindsay, Enterline and Joyce (2008) is explained by the incomplete interpretation of the effects over time (Licht 2011). After correct interpretation, joining on both government and rebel side extends war duration in the Correlates of War sample. Collier, Hoeffler and Soderbom (2004), on the other hand, derive their conclusion due to the assumption that the effect of external support inflates rather than decays over time (Regan 2010).[18] Therefore, empirical models yield that external support prolongs civil conflict.

Unlike other studies, Gent (2008) accounts for the supply of support (see the discussion further below) and demonstrates that military reinforcement of the government has no effect on the duration until government's victory, yet support to the rebels shortens conflict; Gent argues that this finding supports his game-theoretical exercise, because it conforms to the notion that support is given when the recipient cannot win without it. The author makes the case further that, in terms of reinforcing the government side, this means that the

[16] Balch-Lindsay, Enterline and Joyce (2008) use the Correlates of War (COW) sample of civil wars and military interveners.

[17] Collier, Hoeffler and Soderbom (2004) use Regan's (2000, 2002) sample of civil wars and interveners.

[18] Modeling the effect of intervention as a decaying effect over time has been standard in the literature.

government was losing without external support, and in terms of bolstering rebels, it means that the rebels were close to winning. Gent's (2008) contribution is twofold. First, the author develops a formal model of the supply side of external support. Second, the author discusses his empirical findings about war duraion in terms of a selection effect of reinforcement and the process of fighting, although he does not account for the selection effect empirically.

While significant transfers of weapons and equipment are commonplace in interstate wars,[19] there has been a lack of quantitative research on the effect of such transfers on interstate war duration. The effect of multiple states participating in war as opposed to a two-side war duration has been shown to reduce interstate wars (Bennett and Stam 1996). However, participation of multiple states in an interstate war should not be viewed as 100% analogous to sending troops to support one of the sides in a civil war, as some of the time multiple states are original belligerents as opposed to joiners. Since I assume that the underlying logic of conflict is analogous in both intra- and inter-state wars, I expect my model's predictions to apply to both "types" of conflict as long as I account for the characteristics that are more often observed in some conflicts rather than others.[20]

The supply of external support: Based on Hirshleifer (2000), Gent (2008) develops a game theoretic model of the supply of external support in civil wars. Assuming that third parties want to make a difference when they reinforce belligerents, Gent derives that outsiders invest only when their reinforcements have the greatest marginal impact, i.e., "reinforcement is most welcome when it reverses a force disparity from slight inferiority to slight superiority

[19] E.g., consider the US supply of weapons and equipment to Israel during the Yom Kippur War in 1973, or the USSR's weapons supply to the Arab states on the opposing side.

[20] See subsection 1.2 for a more detailed explanation of why I consider the logic of intra- and inter-state wars to be analogous.

(Hirshleifer 2000, p. 776). Employing a version of Regan's data, Gent finds that the stronger the rebels, the more likely the government is to receive outside support.[21] This finding is consistent with the result that Huth (1998) establishes for the sample of interstate crises: the stronger the target the less likely it is to receive reinforcement. The authors imply that by maximizing the recipient's probability of victory, the desired policy outcome is more likely to be achieved. I incorporate the pursuit of third parties to establish a specific policy outcome through reinforcement explicitly in the supply decisions of the extended model.

In the interstate conflict literature, the supply side of joining has been the central question of research on third party influence and war. The most extensive contributions to this literature conform to the framework of opportunity and willingness.[22] By "opportunity" scholars have implied material (i.e., possession of resources) and circumstantial (e.g., geographic proximity) ability to influence the outcome of the conflict. By "willingness" scholars have referred to regime similarity, alliance obligations, rivalry, etc. In other words, in the interstate conflict literature, the focus has been not only on third parties' ability to make a difference (as it has been central to arguing why third parties intervene in civil wars), but also on the source of preferences for intervention. In sum, both inter- and intra-state conflict research have identified a third party's material ability to influence the outcome of a conflict as the main driver of the provision of external support.[23] In the model that follows, I focus

---

[21] Gent (2008) also finds that stronger rebels are more likely to receive outside support, which (the author argues) implies that the rebels are stronger compared to other rebel groups, but not stronger than the government. I revisit this issue in chapter 5.

[22] See (Joyce, Ghosn and Bayer 2013, pp. 4-15) for a detailed literature review on this line of research. Also, see the review of extant literature on this topic by (Melin and Koch 2010, pp. 3-6).

[23] I should note, however, that although some studies, e.g., Melin and Koch (2010) do not test for the effect of third-party capabilities on joining explicitly, they still support the opportunity argument by demonstrating that one of their strongest results is that major powers are much more likely to join than minor powers, i.e., stronger states have more opportunity to join.

on the opportunity component of a third party's decision making and assume that third parties' willingness to join, or preference to provide external support, is given exogenously.

In sum, research on both inter- and intra-state conflict has demonstrated that power mediation delivers peace. It also appears that the current state of the literature established that power mediation generates unstable peace. Although research on the supply of power mediation in conflict is scarce, it appears that scholars agree on the general decision making process that potential mediators undertake: they are more likely to mediate as the benefits of a potential settlement increase and they are less likely to mediate as the costs of mediation effort rise. With respect to external support, military reinforcements seem to have an inconsistent effect on duration in civil conflicts and are likely to shorten interstate conflicts.[24] Scholars of both inter- and intrastate conflict agree that supporters are most likely to take action when the marginal impact of their support is the greatest.

4.1.3 Modeling external support and power mediation

Recap of the decision-making within a fighting dyad: Recall the sequence of events in the benchmark model. In the case either side happens to control more benefits π than what is proportional to his/her military power, a conflict of interests arises. Side A has the power to propose the terms and B rejects or accepts offers. A tries to assess whether B is a weak, a moderately strong, or a strong type. Sometimes A underestimates B's type and the sides fight a battle each time period until A makes a generous enough offer to satisfy B. If A believes that B is strong, A makes a safe (albeit sometimes too generous) offer $\underline{x}_t$, which

[24] The finding about the duration of interstate wars concerns multi-party war participation which may or may not include external support, so this conclusion needs to be viewed with that caveat.

B always accepts. The conflict of interests is resolved. Given that A's initial beliefs have allowed the sides to avoid conflict, A's beliefs stay unchanged.

If A believes that B is moderately strong, she proposes to give up $\hat{x}_t$, which satisfies a weak or a moderately strong B, but is rejected by a strong B. If $\hat{x}_t$ is rejected, the sides fight a battle to redistribute the good according to the military outcome. A's beliefs are updated to reflect the revelation that B is strong. Even if the conflict of interests arises again (in the case when A won a battle and the new distribution of the good still does not reflect that B is strong), the sides will not fight another battle in $t+1$, as A will give up $\underline{x}_t$ to satisfy strong B, as A has already learned that B is strong. If B accepts $\hat{x}_t$, A needs to reevaluate whether given the new information that B is not strong, A continues to believe that B is moderately strong or A believes now that B is weak. If A believes that B is weak, then A risks war by offering $\overline{x}_t$, which B rejects if he is moderately strong and the sides fight a battle.

If A believes B is a weak type, then she makes an offer of $\overline{x}_t$. B will accept such offer if B is in fact weak. However, if B happens to be moderately strong or strong, then B rejects and the sides fight a battle. After the battle, A's beliefs about B's strength are updated based on whether A won or lost. In $t+1$, A either believes that B is strong or moderately strong. If A believes the former then A makes the most generous offer of $\underline{x}_t$ that any B accepts. If A risks another battle by offering $\hat{x}_t$, B's types is revealed fully: if B rejects, B is strong, and the sides fight the second battle. If B accepts, then B is moderately strong, and the conflict is resolved.

Three types of B imply that the maximum of two battles may take place in the null model. The logic of the null model illustrates the learning process of fighting that is generally accepted in the bargaining literature as the mechanism to solve the uncertainty-driven conflicts. Given that no shocks to capabilities occur in the null model, the ceasefires are very stable and never experience another war again. Introducing third parties changes this dynamic.

Recall that A sets the size of x_t using one of five inequalities described in chapter 2. Both the inequalities describing A's decision-making process with respect to determining the initial value of x in period t (the purpose of which is to determine the initial offer, 2.1, 2.2, 2.3) and the inequalities describing A's subsequent decisions to determine the value of x in period $t+1$ (the purpose of which is either to determine the follow-up offer or the offer after the sides have fought their first battle, 2.4, 2.5) have common features. First, the value of the right-hand side increases as the sum of costs that the sides pay when they fight grows. Second, the right-hand side increases when the difference between A's probability of defeating B when B is relatively weaker and when B is relatively stronger shrinks. In other words, the probability of A proposing the terms that ensure peace grows as fighting becomes more costly and as A can assess with more clarity whether B is a relatively stronger or a relatively weaker type.

In the extended model, third parties may choose to pressure the belligerents into ending war, or support either side of the conflict, or stay out. Each third party is denoted as I_i, where i is the third party's numeric identifier. With respect to any given dyad, each third party has a preferred policy outcome (i.e., I's preferred division of the disputed good) denoted as π', utility for side A winning a battle is denoted as $u_{I_{iA}}$, utility for side B's victory is indicated as $u_{I_{iB}}$, and utility for the immediate end of violence donated as $u_{I_{i\Pi}}$. Recall that studies on the occurrences of mediation and support argue that third parties weigh the benefits they may derive from extending their influence against paying the cost of taking action (support and mediation require resources). In other words, third parties evaluate the utility they may derive from the sought outcome weighted by the probability of its occurrence, also subtracting the cost of achieving the sought outcome. Subsection 4.1.2 reviews the literature on where third-party preferences come from: geographic proximity, shared history, and alliance ties are cited as major drivers of “willingness” to mediate or reinforce. A third party’s capability or major power status is referred to as a main proxy for a mediator’s ability to impose a ceasefire. Similarly, the lack of preponderance of power within the fighting dyad is viewed as increased likelihood that a third party’s support will make a difference. I consider the “willingness” component of third parties’ expected utility functions as exogenously given utilities. In contrast, I focus on the “opportunity” component of third party decision making as endogenous (to the model) quantities: the likelihood of the sought outcome and the costs of attaining it.

Supply of support: First, a third party that has a preference for support should check two conditions: policy alignment and the recipient’s need for support. Determining policy alignment means assessing the gap between I_i’s preferred policy outcome π' and the current distribution of benefits π. Recall that π denotes the amount of benefits in the dyad controlled by side A, and $1-\pi$ indicates B’s benefits in the dyad. A third party I has a preference for reinforcing A and will consider reinforcing A (i.e., will carry out the calculation of her/his expected utility from such action) if three conditions hold:

- the utility that I acquires from A’s victory is greater than that from B’s victory, $u_{I_{iA}} > u_{I_{iB}}$;

- a third party prefers A's victory to an immediate end of violence, $u_{I_{iA}} > u_{I_\Pi}$;
- the current division of the disputed good is distributed in a way that diverges from I's preferred policy outcome, $\pi \leq \pi'$.

With respect to the third condition, if the current status quo is greater than π' the distribution of benefits that A receives is either equal to or exceeds I_i's preferred outcome, therefore I_i is not motivated to influence the policy outcomes for the dyad in question. If, however, $\pi < \pi'$, then the distribution of benefits that A receives is less than I_i's preferred policy outcome and I_i is motivated to influence who gets what in the dyad in question (provided two other conditions are true as well).

Before I_i is able to evaluate the expected utility from reinforcing A, I_i determines the optimal value of his/her investment m_{I_i} that I_i has to contribute to increase A's probability of military victory. Recall that the probability of A's winning is $p_t = \dfrac{m_A}{m_A + m_B + b}$. Since I_i does not know the value of B's private capabilities, I assume that I_i can see the same observable signals that A uses to form her beliefs about B's private capabilities b.[25] Therefore, I assume that each third party that considers investing in A, uses A's decision about the value of x_t to guess B's type.[26] If I_i provides some support $m_{I_{iA}}$ to A, then A's probability of winning is $p_t = \dfrac{m_A + m_{I_{iA}}}{m_A + m_{I_{iA}} + m_B + b}$. Battles are won probabilistically and resources in the model are distributed in accordance with relative power. This means that a third party that is motivated by a certain policy outcome can attempt to bring A's probability of winning p_t as close to the desired policy outcome, π', as possible. Solving for $m_{I_{iA}}$ if p_t equals π' yields:

[25] In reality, beliefs are based on prior behavior. In the model, beliefs are exogenous, and are assumed to be based on B's prior behavior.

[26] See the extended model's outline that details the stages of A's decision to propose x_t: first, A makes a decision as if third parties are not present, then third parties decide whether to invest or pressure for peace, then A reexamines the size of x_t. To calculate the size of m_I, I uses the initial value of x_t, which reflects A's assessment of B's type: $\underline{x}_t$ implies that A thinks that B is strong, $\hat{x}_t$ - moderately strong, and $\overline{x}_t$ weak.

$$m_{I_{iA}} = \frac{\pi' m_B + \pi' b - (1 - \pi') m_A}{(1 - \pi')} \tag{4.1}$$

Therefore, a third party's participation depends on whether I_i is able to contribute enough resources such that the preferred policy outcome could be achieved by the recipient through bargaining.

The size of $m_{I_{iA}}$ is the cost of external support that I_i has to pay to reinforce A. The cost grows as B's relative advantage increases and it becomes negligible if A dominates B. For instance, if A controls 10% of the dyad's capabilities, then to bring A's chance of winning close to 51% (which implies that I_i wants A to control .51 of the dyad's benefits), I_i would need to invest capability of 0.92. Given that every participant in the system controls at most a capability of 1.0, I_i would need to contribute almost all of her/his resources. To prevent such occurrences, I assume that no investment may exceed 30% of I_i's total capabilities, m_{IT}. In contrast, if A dominates B by controlling 90% of capabilities in the dyad, the determined size of $m_{I_{iA}}$ would be $-.79$. Given that resources cannot be negative, I assume that I_i only considers the value of his/her expected utility calculation if $m_{I_{iA}}$ is positive to reflect the empirical reality that reinforcement is mostly provided to the belligerents who need it to win and achieve the desired policy. I should note, however, that the assumption to avoid negligible amounts of reinforcement toward successful belligerents eliminates the phenomenon of band wagoning in my model.[27]

[27] See Joyce, Ghosn and Bayer (2013) for a discussion of band wagoning in international crises. The motivation for band wagoning, the authors argue, is extraction of future benefits from a powerful target or initiator in exchange for providing the belligerent some extra legitimacy by the act of joining the conflict. The motivation for support in my model is third party's ability to make a difference in the military outcome of war, therefore, band wagoning is beyond the scope of this model.

If the three conditions for I_i's preference are true and if m_I is positive and is less than 30% of I_i's capabilities, then I_i calculates his/her expected utility for external support, and if this utility is greater than zero, I_i provides m_{I_iA} to A. I assume that the greater the disparity between π and π', the more motivated I_i is to reinforce A. Furthermore, the intensity of preference for A's victory, u_{I_A}, over side B's victory, u_{I_B}, is also proportional to I_i's motivation to build up A. Finally, the motivation for investment should exceed the cost of investment, m_{I_i}. I assume away some cost that is associated with the logistics of providing support. The expected utility for supporting A is shown in equation 4.2.

Similarly, if I_i favors side B, I_i first determines whether $u_{I_{iA}} < u_{I_{iB}}$, and $u_{I_{iB}} > u_{I_\Pi}$ and $(1-\pi) < 1-\pi'$. If so, the distribution of benefits that B receives is less than I_i's preferred policy outcome and I_i is motivated to influence who gets what in the dyad in question. Then, I_i determines whether B needs his/her support by calculating $m_{I_{iB}} = \dfrac{\pi' m_A - (1-\pi')m_B - (1-\pi')b}{\pi'}$. Again, I_i uses the value of b that is inferred by A. If $m_{I_{iB}} > 0$ and $m_{I_{iB}} < .3m_B$, then I_i calculates the expected utility for providing external support to B, as shown in equation 4.2.

$$EU_I(ES) = \begin{cases} (u_{I_A} - u_{I_B})(\pi' - \pi) - m_{I_A}, & \text{if } u_{I_A} > u_{I_B}, u_{I_A} > u_{I_\Pi}, \pi \leq \pi', m_{I_A} \in (0, 0.3m_{IT}] \\ \\ (u_{I_B} - u_{I_A})(\pi - \pi') - m_{I_B}, & \text{if } u_{I_A} < u_{I_B}, u_{I_B} > u_{I_\Pi}, \pi \geq \pi', m_{I_B} \in (0, 0.3m_{IT}] \end{cases} \tag{4.2}$$

Supply of power mediation: If I_i derives greater utility from ending violence than either side's victory, i.e. $u_{I_{i\Pi}} > u_{I_{iA}}$ and $u_{I_{i\Pi}} > u_{I_{iB}}$ then I assume that a third party is potentially motivated to pressure the belligerents to stop fighting. To do so, a third party needs to be motivated, the outcome of such pressure should not vastly contradict a third party's preferred policy outcome, and a third party needs to be powerful enough to carry out such pressure, i.e. threats of military intervention, or withdrawal of resources, or sanctions need to be believable. The utility function of support (as summarized in equation4.4) consists of two major terms: the cost of pressure are subtracted from the utility derived from pressure. With respect to the utility term, the nominator is proportional to the intensity with which a third party prefers end of violence over supporting either side, while the denominator gets smaller proportionally to the absolute difference between I_i's preferred policy outcome π' and the current division of the disputed good π. Thus, the more I_i prefers end of violence over supporting either side and the more closely the current status quos reflects I_i's preferred policy, the more likely a third party is to inflate the costs of fighting for the belligerents A and B by using power mediation.

The cost term consists of another fraction. The costs are greater as the total capabilities of the dyad increase (the nominator) to reflect that it is more costly to impose a settlement on major powers as opposed to minor powers. In contrast, the cost term shrinks as the proportion of the capabilities in the dyad controlled by I_i grows, therefore the denominator contains the proportion $\frac{m_{IT}}{m_A + m_B + b + m_{IT}}$, as shown in equation 4.4. In the presented version of the model, if a third-party agent exerts pressure on the belligerents, that agent pays the cost of power mediation m_{IPM} as shown in the equation 4.3.[28]

[28] It is, however, probable that the expected amount of resources to be able to generate believable threats (the shown cost term) is a greater amount of resources than what power mediators do pay when they pressure the belligerents to stop fighting. The current version of the model most likely exaggerates how costly pressure for peace is; I will revisit this cost term in future versions of the model.

$$m_{IPM} = \frac{m_A + m_B}{\frac{m_{IT}}{m_A + m_B + b + m_{IT}}} \tag{4.3}$$

If $EU_{I_i}(PM) > 0$, then c_A and c_B in period t are reset to $c_{At} = c_{A(t-1)} + \frac{m_{IT}}{m_A + m_B + b + m_{IT}}$ and $c_{Bt} = c_{B(t-1)} + \frac{m_I}{m_A + m_B + m_{IT}}$. This adjustment in the size of belligerents' costs of fighting reflects the amount of pressure that a power mediator exerts on the belligerents. The greater the power of I_i compared to A and B, the easier it is for I_i to make threats believable and exert prohibitive costs of fighting on the belligerents. A one time provision of power mediation keeps the belligerents' costs of fighting inflated for either one or three time periods. The effectiveness of power mediation to keep a dyad from fighting is the direct result of such an assumption. The logic of an imposed ceasefire crumbling is the same regardless of the assumption: once the impact of power mediation ends, the artificially inflated costs of fighting deflate to their original values, which leads to the participants' dissatisfaction with the status quo. See the discussion of illustrative scenarios further below that explains the deviations from this dynamic due to the influence of other power mediators and support-providing third parties.

$$EU_T(PM) = \begin{cases} \frac{u_{I_\Pi} - u_{I_A}}{|\pi' - \pi|} - m_{IPM}, & \text{if } u_{I_\Pi} > u_{I_A}, u_{I_\Pi} > u_{I_B}, u_{I_A} > u_{I_B}, m_{IPM} \in (0, 0.3m_{IT}] \\ \frac{u_{I_\Pi} - u_{I_B}}{|\pi' - \pi|} - m_{IPM}, & \text{if } u_{I_\Pi} > u_{I_B}, u_{I_\Pi} > u_{I_A}, u_{I_B} > u_{I_A}, m_{IPM} \in (0, 0.3m_{IT}] \end{cases} \tag{4.4}$$

4.2 Pseudocode of the Extended Model

The sequence of steps programmed to reflect the discussion above follow. Consider Tables 4.1-4.2 for the description and summary of the initialization parameters and a list of actor characteristics/quantities of interest introduced in this chapter. Please consult Tables 2.1-2.2 for the parameters introduced in the second chapter.

4.2.1 Big picture outline

The major steps of the simulation remain the same compared to the null model. The distinctions from the null model are reflected in additional steps included in loops 1, 3 and 4 (shown in bold), which include third parties' decision-making processes about whether to provide external support or power mediation and the belligerents incorporating third-party influence into their bargaining interaction.

1. Initialize the world, set up the agents and the dyads. For each dyad, B is assigned a type (weak, moderately strong, or strong) randomly, **set up third parties for each dyad.**
2. For each dyad, check whether either side is dissatisfied with the status quo. If both are satisfied, the dyad stays at peace. If either side is dissatisfied, the dyad's status is updated to having a conflict of interests;
3. For each dyad that has a conflict of interests: A determines which terms to propose based on her beliefs and the difference between defeating two different types of B;
    - **For each dyad that has a conflict of interests: for each third party in relation to every given dyad: calculate the expected utilities from providing external support or power mediation, when applicable. Update the capabilities and cost terms for each dyad if third parties have exerted influence;**
    - **Given that the cost and capability terms may have changed due to third-party influence, re-check whether either side continues to be dissatisfied with the status quo. If both are satisfied, the dyad is reset as**

peaceful without any transfers of resources from A to B. If either side continues to be dissatisfied, the dyad's status continues to indicate a conflict of interests;

- **For each dyad that continues to have a conflict of interests: A redetermines which terms to propose based on her beliefs and the difference between defeating two different types of B, which now (for some dyads) incorporate influenced by third parties costs of fighting and relative capabilities:**
  - If A proposes the terms that all three types accept, the dyad's status is updated to peaceful.
  - If A proposes the terms that only a moderately strong and a weak types accept:
    (a) If B is strong, B rejects. A's beliefs are reset. Uncertainty solved. The dyad's status is updated to armed conflict.
    (b) If B is moderately strong or weak, B accepts. Uncertainty remains. Is A's belief that B is weak such that A risks war?
      i. If no, the terms are set to reflect that B is moderately strong; the dyad's status is updated to peaceful.
      ii. If yes, A proposes the terms that only weak B would accept.
        A. If B accepts, B is weak, the beliefs and terms are reset to reflect this; the dyad's status is updated to peaceful.
        B. If B rejects, B is moderately strong, the beliefs are reset to reflect this, uncertainty is solved; the dyad's status is updated to armed conflict.
  - If A proposes the terms that only a weak type would accept:
    (a) If B is weak, B accepts and the dyad's status is set to peaceful.
    (b) If B is moderately strong or strong, B rejects, uncertainty is not solved; the dyad's status is reset as armed conflict.

4. For each dyad that is at armed conflict:

- If uncertainty is solved. The sides fight one battle.
  (a) If B wins, the terms are set to reflect that B is of a stronger type.
  (b) If A wins, the terms are set to the last offer, however given A's updated beliefs, the dyad will not escalate to armed conflict next period.
- If uncertainty is not solved. Uncertainty is over whether B is moderately strong or strong. The sides fight a battle. If A wins, her belief that B is moderately strong increases. If B wins, A's belief that B is moderately strong decreases.
  - **For each dyad that has fought a battle and in which uncertainty is not solved: for each third party in relation to every such dyad: calculate the expected utilities from providing external support or power mediation, when applicable. Update the capabilities and cost terms for each dyad if third parties have exerted influence;**

- A **compares her updated belief about** B **being moderately strong conditional on the outcome of the first battle to** A**'s utility of proposing the terms as if** B **is moderately strong, as opposed to proposing the terms as if** B **is strong. For some dyads these values are influenced by third parties, as they involve cost terms and probabilities of military victory.**
  (a) If A estimates thats B is moderately strong, A proposes the terms that only a moderately strong B accepts.
    i. If B is moderately strong, he accepts, the dyad status is peaceful.
    ii. If B is strong, he rejects, A's beliefs are updated to reflect that; uncertainty is solved; the dyad's status stays at armed conflict.
  (b) If A estimates now that B is strong, A proposes the terms that any type of B accepts. The terms are updated.

5. Go to step 2.

4.2.2 Detailed pseudocode

The following sequence of steps is written in the Java programming language with the use of RePast libraries to reflect the discussion above.

I. **Do these steps once:**

1. Create a `numstates` number of political actors in the "world" (minimum of 2), each characterized by randomly generated $m_i \in [0, 1]$.
2. All actors are paired up in non-directional dyads and are assigned randomly either the role of the potential defender A, or the role of the potential challenger B.
3. All dyads are located on the `ArrayOfPeacefulDyads`.
4. Each dyad is assigned a randomly generated value π that reflects the division of the disputed good in the dyad. π is the proportion of the good controlled by A.
5. Each dyad is assigned two randomly generated beliefs for A: ϕ_w and ϕ_m, such that $\phi_w + \phi_m \leq 1$.
6. Each dyad is assigned three randomly generated values of b for B, such that $\underline{b} < \hat{b} < \overline{b}$ AND $b \leq (0.2 \times m_B)$.
7. Set `termsDeterminedInPreviousRound`=FALSE for each dyad.
8. **Set up third parties for every dyad:**
   - Generate a `thirdPartyArray` for each dyad that contains all agents in the system except this dyad's belligerents; - For each dyad: for each third party: generate randomly $u_{I_A}, u_{I_B}, u_{I_\Pi}$ to describe this party's preferred outcome in each dyad.

To ensure that events take place sequentially, execute for-loop 1 and for-loop 2 if the tick is even, execute for-loop 3 if the tick is odd.

II. **For-loop 1:** for each dyad located on the `ArrayOfPeacefulDyads`, each tick follows this sequence:

1. Determine whether a conflict of interests exists. In each dyad, B checks whether $\pi > p_t + c_B$ (B knows which type he is, so the choice of p_t is obvious for B) and A checks whether $\pi < p_t - c_A$. With respect to the choice of p_t, A first calculates, $\phi_w^{k_1} = \frac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B} - \phi_m$, $\hat{\phi}_w^{k_2} = \frac{c_A + c_B}{\overline{p}_t - \underline{p}_t + c_A + c_B} - \phi_m \frac{\hat{p}_t - \underline{p}_t}{\overline{p}_t - \underline{p}_t + c_A + c_B}$, $\hat{\phi}_w^{k_3} = \phi_m \frac{c_A + c_B}{\overline{p}_t - \underline{p}_t}$. When $\phi_w \leq min\{\hat{\phi}_w^{k_1}, \hat{\phi}_w^{k_2}\}$, A uses $\underline{p}_t$; when $\phi_w > max\{\hat{\phi}_w^{k_2}, \hat{\phi}_w^{k_3}\}$, A uses $\overline{p}_t$; when $\hat{\phi}_w^{k_1} < \phi_w < \hat{\phi}_w^{k_3}$, A uses $\hat{p}_t$.
   i. If either condition holds, the dyad is added to the `ArrayOfConflictsOfInterest`.

III. **For-loop 2:** for each dyad located on the `ArrayOfConflictsOfInterest`, each tick follows this sequence:

1. If `termsDeterminedInPreviousRound`=FALSE, determine x_t. When $\phi_w \leq min\{\hat{\phi}_w^{k_1}, \hat{\phi}_w^{k_2}\}$, A sets $x_t = \underline{x}_t$, when $\phi_w > max\{\hat{\phi}_w^{k_2}, \hat{\phi}_w^{k_3}\}$, A sets $x_t = \overline{x}_t$, when $\hat{\phi}_w^{k_1} < \phi_w < \hat{\phi}_w^{k_3}$, A sets $x_t = \hat{x}_t$. Record the original x_t to file.
2. For each third party selected randomly from a population of third parties in relation to this dyad: the selected agent I_i determines whether s/he is interested in support or power mediation. The influenced capability and cost terms are changed immediately, such that the next third party that considers exerting influence takes the contributed support/mediation into account.
   - If $u_{I_A} > u_{I_B}$ AND $u_{I_A} > u_{I_\Pi}$ AND $\pi \leq \pi', m_{I_A} \in (0, 0.3m_{IT}]$, then I_i calculates $EU(ES) = (u_{I_A} - u_{I_B})(\pi' - \pi) - m_{I_A}$. If $EU(ES) > 0$, then reset $m_A = m_A + m_{I_A}$.
   - If $u_{I_A} < u_{I_B}, u_{I_B} > u_{I_\Pi}$ AND $\pi \geq \pi'$ AND $m_{I_B} \in (0, 0.3m_{IT}]$, then I_i calculates $EU(ES) = (u_{I_B} - u_{I_A})(\pi - \pi') - m_{I_B}$. If $EU(ES) > 0$, then reset $m_B = m_B + m_{I_B}$.
   - If $u_{I_\Pi} > u_{I_A}$ AND $u_{I_\Pi} > u_{I_B}$ AND $u_{I_A} > u_{I_B}$, then I_i calculates $EU(PM) = \frac{u_{I_\Pi} - u_{I_A}}{|\pi' - \pi|} - \frac{m_A + m_B}{\frac{m_{IT}}{m_A + m_B + b + m_{IT}}}$, if $EU(PM) > 0$, then reset $c_A = c_A + \frac{m_{IT}}{m_A + m_B + b + m_{IT}}$ AND $c_B = c_B + \frac{m_{IT}}{m_A + m_B + b + m_{IT}}$
   - If $ifu_{I_\Pi} > u_{I_B}$ AND $u_{I_\Pi} > u_{I_A}$ AND $u_{I_B} > u_{I_A}$, then I_i calculates $EU(PM) = \frac{u_{I_\Pi} - u_{I_B}}{|\pi' - \pi|} - \frac{m_A + m_B}{\frac{m_{IT}}{m_A + m_B + b + m_{IT}}}$, if $EU(PM) > 0$, then reset $c_A = c_A + \frac{m_{IT}}{m_A + m_B + b + m_{IT}}$ AND $c_B = c_B + \frac{m_{IT}}{m_A + m_B + b + m_{IT}}$

3. Given that some dyads have new values for p_t and c, re-determine whether a conflict of interests exists. In each dyad, B again checks whether $\pi > p_t + c_B$ (B knows which type he is, so the choice of p_t is obvious for B) and A checks whether $\pi < p_t - c_A$. With respect to the choice of p_t, A first calculates, $\hat{\phi}_w^{k_1} = \dfrac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B} - \phi_m$, $\hat{\phi}_w^{k_2} = \dfrac{c_A + c_B}{\overline{p}_t - \underline{p}_t + c_A + c_B} - \phi_m \dfrac{\hat{p}_t - \underline{p}_t}{\overline{p}_t - \underline{p}_t + c_A + c_B}$, $\hat{\phi}_w^{k_3} = \phi_m \dfrac{c_A + c_B}{\overline{p}_t - \underline{p}_t}$. When $\phi_w \leq min\{\hat{\phi}_w^{k_1}, \hat{\phi}_w^{k_2}\}$, A uses $\underline{p}_t$; when $\phi_w > max\{\hat{\phi}_w^{k_2}, \hat{\phi}_w^{k_3}\}$, A uses $\overline{p}_t$; when $\hat{\phi}_w^{k_1} < \phi_w < \hat{\phi}_w^{k_3}$, A uses $\hat{p}_t$.
   i. If either condition holds, the dyad is kept on the `ArrayOfConflictsOfInterest`.
   ii. If neither condition holds any longer, the dyad is moved back to the `ArrayOfPeacefulDyads`.
4. For those dyads that have remained on the `ArrayOfConflictsOfInterest`, re-determine x_t. When $\phi_w \leq min\{\hat{\phi}_w^{k_1}, \hat{\phi}_w^{k_2}\}$, A sets $x_t = \underline{x}_t$, when $\phi_w > max\{\hat{\phi}_w^{k_2}, \hat{\phi}_w^{k_3}\}$, A sets $x_t = \overline{x}_t$, when $\hat{\phi}_w^{k_1} < \phi_w < \hat{\phi}_w^{k_3}$, A sets $x_t = \hat{x}_t$.
   i. If A proposes terms $\underline{x}_t$:
      A. Every type of B accepts.
         - Update $\pi = \underline{p}_t + c_B$. Move the dyad to the `ArrayOfPeacefulDyads`.
   ii. If A proposes terms $\hat{x}_t$:
      A. If B is strong, he rejects.
         - Move the dyad to the `ArrayOfArmedConflicts`. Set boolean `uncertaintySolved` as true. Since only a strong type rejects, update A's beliefs as: $\phi_w = 0.05, \phi_m = 0.05$.
      B. If B is weak or moderately strong, he accepts.
         - if $\phi_w \leq \dfrac{c_A + c_B}{\overline{p}_t - \hat{p}_t + c_A + c_B}$:

           A keeps $\hat{x}_t$. Move the dyad to the `ArrayOfPeacefulDyads`. A's beliefs are updated to reflect that A knows that B is not strong, so the belief that B is strong is reduced to $(1 - \phi_w - \phi_m) = 0.05$, the rest of the difference is distributed between ϕ_m and ϕ_w equally.
         - if $\phi_w > \dfrac{c_A + c_B}{\overline{p}_t - \hat{p}_t + c_A + c_B}$:

           The conflict of interest persists, as A now thinks it is possible to keep $\overline{x}$ in the following round. Keep the dyad on this array and set `termsDeterminedInPreviousRound` as TRUE.
   iii. If A proposes terms $\overline{x}_t$ and `termsDeterminedInPreviousRound` is FALSE:
      A. If B is weak, he accepts.
         - Update $\pi = \overline{p}_t + c_B$. Update A's beliefs $\phi_w = 0.99$ and $\phi_m = 0.05$ to reflect that A knows that B is weak. Move the dyad to the `ArrayOfPeacefulDyads`.
      B. If B is moderately strong or strong, he rejects.
         - Move the dyad to the `ArrayOfArmedConflicts`. Set boolean `uncertaintySolved` as false.

5. If `termsDeterminedInPreviousRound`=true, do:
   i. A proposes terms $\overline{x}_t$ and `termsDeterminedInPreviousRound` is TRUE
      A. If B is weak, he accepts.
         • Update $\pi = \overline{p}_t + c_B$. Update A's beliefs $\phi_w = 0.99$ and $\phi_m = 0.05$ to reflect that A knows that B is weak. Move the dyad to the `ArrayOfPeacefulDyads`.
      B. If B is moderately strong, he rejects.
         • Move the dyad to the `ArrayOfArmedConflicts`. Set boolean `uncertaintySolved` as true. Since only moderately strong B would reject that offer, A knows B's type now, update beliefs such that $\phi_w = 0.05, \phi_m = 0.99$

IV. **For-loop 3:** for each dyad located on the `ArrayOfArmedConflicts`, each tick do:

1. The sides fight a battle. Generate a random value p' that is drawn from a uniform distribution, $p' \in [0, 1]$.
2. If `uncertaintySolved` is TRUE and B rejected $\hat{x}_{t-1}$:
   i. If $p' \leq \hat{p}_t$, A wins the battle, A keeps $\hat{x}_t$; update $\pi = \hat{p}_t + c_B$; move the dyad to the `ArrayOfPeacefulDyads`.

   ii. If $p' > \hat{p}_t$, A loses, she keeps $\underline{x}_t$; update $\pi = \underline{p}_t + c_B$; move the dyad to the `ArrayOfPeacefulDyads`.
3. If `uncertaintySolved` is TRUE and B accepted $\hat{x}_{t-1}$, but rejected $\overline{x}_{t-1}$:
   i. If $p' \leq \hat{p}_t$, A wins the battle, A keeps $\overline{x}_t$; update $\pi = \overline{p}_t + c_B$; move the dyad to the `ArrayOfPeacefulDyads`.

   ii. If $p' > \hat{p}_t$, A loses, she keeps $\hat{x}_t$; update $\pi = \hat{p}_t + c_B$; move the dyad to the `ArrayOfPeacefulDyads`.
4. If `uncertaintySolved` is FALSE (B rejected $\overline{x}_{t-1}$):
   i. If $p' \leq \hat{p}_t$, A wins a battle. A calculates updated belief about the value of ϕ_m: calculates $\phi_m^w = \dfrac{\hat{p}_t \phi_m}{\hat{p}_t \phi_m + \underline{p}_t (1 - \phi_m)}$.
   ii. For each third party selected randomly from a population of third parties in relation to this dyad: the selected agent I_i determines whether s/he is interested in support or power mediation. The influenced terms are changed immediately, such that the following third party takes the contributed support/mediation into account.
      • If $u_{I_A} > u_{I_B}$ AND $u_{I_A} > u_{I_\Pi}$ AND $\pi \leq \pi', m_{I_A} \in (0, 0.3 m_{IT}]$, then I_i calculates $EU(ES) = (u_{I_A} - u_{I_B})(\pi' - \pi) - m_{I_A}$. If $EU(ES) > 0$, then reset $m_A = m_A + m_{I_A}$.
      • If $u_{I_A} < u_{I_B}, u_{I_B} > u_{I_\Pi}$ AND $\pi \geq \pi'$ AND $m_{I_B} \in (0, 0.3 m_{IT}]$, then I_i calculates $EU(ES) = (u_{I_B} - u_{I_A})(\pi - \pi') - m_{I_B}$. If $EU(ES) > 0$, then reset $m_B = m_B + m_{I_B}$.

- If $u_{I_\Pi} > u_{I_A}$ AND $u_{I_\Pi} > u_{I_B}$ AND $u_{I_A} > u_{I_B}$, then I_i calculates $EU(PM) = \dfrac{u_{I_\Pi} - u_{I_A}}{|\pi' - \pi|} - \dfrac{m_A + m_B}{\dfrac{m_{IT}}{m_A + m_B + b + m_{IT}}}$, if $EU(PM) > 0$, then reset $c_A = c_A + \dfrac{m_{IT}}{m_A + m_B + b + m_{IT}}$ AND $c_B = c_B + \dfrac{m_{IT}}{m_A + m_B + b + m_{IT}}$
- If $if u_{I_\Pi} > u_{I_B}$ AND $u_{I_\Pi} > u_{I_A}$ AND $u_{I_B} > u_{I_A}$, then I_i calculates $EU(PM) = \dfrac{u_{I_\Pi} - u_{I_B}}{|\pi' - \pi|} - \dfrac{m_A + m_B}{\dfrac{m_{IT}}{m_A + m_B + b + m_{IT}}}$, if $EU(PM) > 0$, then reset $c_A = c_A + \dfrac{m_{IT}}{m_A + m_B + b + m_{IT}}$ AND $c_B = c_B + \dfrac{m_{IT}}{m_A + m_B + b + m_{IT}}$

iii. A then recalculates her probabilities of military victory when B is moderately strong, $\hat{p}_t$, and when B is strong, $\underline{p}_t$ and compares the new belief about ϕ_m to the value that is derived from comparing A's utility of proposing the terms as if B is moderately strong, as opposed to proposing the terms as if B is strong. For some dyads the cost and probabilities of military victory terms are updated due to third party influence:

A. If $\phi_m^w \leq \dfrac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B}$. A sets x_{t+1} equal to $\underline{x}_{t+1}$.
   - B accepts. Update $\pi = \underline{p}_t + c_B$. A's beliefs are updated as $\phi_w = 0.05$ (to reflect that A has found out that B is not weak), and $\phi_m = \phi_m^w$ (to reflect the updated belief ϕ_m after fighting a battle).[29] Move the pair to the `ArrayOfPeacefulDyads`.

B. If $\phi_m^w > \dfrac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B}$. A sets x_{t+1} equal to $\hat{x}_{t+1}$.
   - If B is moderately strong, he accepts, update $\pi = \hat{p}_t + c_b$; A's beliefs are updated to reflect that only a moderately strong type would accept the offer: $\phi_w = 0.05, \phi_m = 0.99$; move the dyad to `ArrayOfPeacefulDyads`.

   - If B is strong, he rejects. A's beliefs are updated to reflect that only a strong type would reject the offer: $\phi_w = 0.05, \phi_m = 0.05$. Set boolean `uncertaintySolved` as TRUE, keep on this array.

iv. If $p' > \hat{p}_t$, B won. A calculates updated belief about the value of ϕ_m: $\phi_m^l = \dfrac{(1 - \hat{p}_t)\phi_m}{(1 - \hat{p}_t)\phi_m + (1 - \underline{p}_t)(1 - \phi_m)}$.

v. For each third party selected randomly from a population of third parties in relation to this dyad: the selected agent I_i determines whether s/he is interested in support or power mediation. The influenced terms are updated immediately, such that the following third party takes the contributed support/mediation into account.

[29] Even though A has won, the most generous offer of $\underline{x}_{t+1}$ reflects that A's belief that B is strong is the highest, therefore, the dyad is extremely unlikely to return to the conflict of interests stage, and if it does, A clearly believes that B is most likely to be strong and therefore will again make an appeasing offer.

- If $u_{I_A} > u_{I_B}$ AND $u_{I_A} > u_{I_\Pi}$ AND $\pi \leq \pi', m_{I_A} \in (0, 0.3m_{IT}]$, then I_i calculates $EU(ES) = (u_{I_A} - u_{I_B})(\pi' - \pi) - m_{I_A}$. If $EU(ES) > 0$, then reset $m_A = m_A + m_{I_A}$.
- If $u_{I_A} < u_{I_B}, u_{I_B} > u_{I_\Pi}$ AND $\pi \geq \pi'$ AND $m_{I_B} \in (0, 0.3m_{IT}]$, then I_i calculates $EU(ES) = (u_{I_B} - u_{I_A})(\pi - \pi') - m_{I_B}$. If $EU(ES) > 0$, then reset $m_B = m_B + m_{I_B}$.
- If $u_{I_\Pi} > u_{I_A}$ AND $u_{I_\Pi} > u_{I_B}$ AND $u_{I_A} > u_{I_B}$, then I_i calculates $EU(PM) = \dfrac{u_{I_\Pi} - u_{I_A}}{|\pi' - \pi|} - \dfrac{\dfrac{m_A + m_B}{m_{IT}}}{m_A + m_B + b + m_{IT}}$, if $EU(PM) > 0$, then reset $c_A = c_A + \dfrac{m_{IT}}{m_A + m_B + b + m_{IT}}$ AND $c_B = c_B + \dfrac{m_{IT}}{m_A + m_B + b + m_{IT}}$
- If $if u_{I_\Pi} > u_{I_B}$ AND $u_{I_\Pi} > u_{I_A}$ AND $u_{I_B} > u_{I_A}$, then I_i calculates $EU(PM) = \dfrac{u_{I_\Pi} - u_{I_B}}{|\pi' - \pi|} - \dfrac{\dfrac{m_A + m_B}{m_{IT}}}{m_A + m_B + b + m_{IT}}$, if $EU(PM) > 0$, then reset $c_A = c_A + \dfrac{m_{IT}}{m_A + m_B + b + m_{IT}}$ AND $c_B = c_B + \dfrac{m_{IT}}{m_A + m_B + b + m_{IT}}$

vi. A then compares the new belief about ϕ_m to the value that is derived from comparing A's utility of proposing the terms as if B is moderately strong, as opposed to proposing the terms as if B is strong, using newly updated $\hat{p}_t$, $\underline{p}_t$, c_A, and c_B:

A. If $\phi^l_m \leq \dfrac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B}$. A sets x_{t+1} equal to $\underline{x}_{t+1}$.

- B accepts. Record $\pi = \underline{p}_t + c_B$. A's beliefs are updated as $\phi_w = 0.05$ (to reflect that A has found out that B is not weak), and $\phi_m = \phi^w_m$ (to reflect the updated belief ϕ_m after fighting a battle).[30] Move the dyad to the `ArrayOfPeacefulDyads`.

B. If $\phi^l_m > \dfrac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B}$. A sets x_{t+1} equal to $\hat{x}_{t+1}$.

- If B is moderately strong, he accepts, update $\pi = \hat{p}_t + c_B$. A's beliefs are updated to reflect that only a moderately strong type would accept this offer: $\phi_w = 0.05, \phi_m = 0.99$. Move the dyad to `ArrayOfPeacefulDyads`.

- If B is strong, he rejects. A's beliefs are updated to reflect that only a strong type would reject the offer: $\phi_w = 0.05, \phi_m = 0.05$. Set boolean `uncertaintySolved` as TRUE. Keep the dyad on this array.

[30] B has won, the most generous offer of $\underline{x}_{t+1}$ reflects that A's belief that B is strong is the highest, therefore, the dyad is extremely unlikely to return to the conflict of interests stage, and if it does, A clearly believes that B is most likely to be strong and therefore will again make an appeasing offer.

4.3 Analysis of Computer Simulations' Numerical Output

4.3.1 Illustrative runs of the extended model

To facilitate comparison with the null model, the default initialization parameter values in the extended model remain the same: "world" system of 200 agents, paired up randomly into 50 dyads (the range of possible initialization parameter values that appear in both versions is shown in Table 2.1 and the initialization parameters specific to the extended version are listed in Table 4.1). The agents are paired up randomly, such that one agent participates as a belligerent in one dyad. With respect to all other dyads, in which a given agent does not participate as a belligerent, this agent may participate as a third party through providing external support or power mediation (given positive expected utility from such actions and provided that certain conditions are met, as outlined in equations 4.2-4.4). Just like in the null model, a dyad may either be at peace, or undergo a conflict of interests, or fight an armed conflict. In contrast to the null model, four interconnected phenomena arise in the extended model:

- The system overall appears to have a higher volume of conflict,
- Conflicts of interests and wars may stop without transfers of resources,
- Conflicts of interests and wars may reignite, i.e., peace spells end,
- A higher proportion of armed conflicts lasts the maximum number of rounds.

Consider panels 1-4 in Figure 4.1, which present the time series of peaceful dyads (shown in green), dyads with conflicts of interest (blue), and fighting dyads (red) over simulation time for various numbers of dyads in the system. Panels (a) through (d) show the systems with 20, 50, 75, and 100 dyads. All simulated "worlds" share a major key feature that sets them apart from the null version of the model: the system never achieves "global peace." For comparison, consider Figure 2.1 that shows analogous time series for the null model, a major takeaway from which is that war is an effective instrument for conflict resolution, as it allows the belligerents to solve the information problem and redivide the good in accordance with relative power in the dyad. In the extended model, war does not appear to be as effective at achieving long-term settlements, as A's beliefs keep being updated throughout the simulation time and the division of benefits is constantly revisited. Consider Figure 4.2 that plots absolute changes in A's beliefs over time in a single simulation run of the extended model (50 dyads are shown). Furthermore, Figure 4.3 charts the absolute changes in the distribution of benefits for 50 dyads. In short, conflicts recur regularly in the extended model. As a result, the volume of conflict is higher in the extended model compared to the null version. Two major reasons explain this emergent property of the extended model: power mediation interrupts conflicts prematurely and transfers of resources from third parties to conflicting dyads undermines the settlements in third parties' dyadic interactions.

First, ceasefires that are imposed by power mediators fail once mediation's effect ends. Power mediation works through inflating the costs of fighting and thus making the sides more likely to settle as increased costs of fighting impel A to propose more generous terms of settlement. A conflict of interests in the null model only ends due to redistribution of benefits in accordance with relative power. In contrast, in the extended model, if a conflict of interests is forced into peace by power mediation, redistribution does not occur and, therefore, the conflict of interests eventually reappears. Thus, the same conflict of interests would generate two conflicts of interests in the extended model, in case power mediation is applied during the conflict of interests stage. Given that it is possible for the same conflict of interests to be pressured into ceasefire multiple times, the difference between the null and the extended model may be even greater. When mediation is applied after the first battle during the armed conflict phase, it (most likely) precipitates a ceasefire such that the sides have a lower chance to remove the information problem completely and, if so, they will reenter a conflict of interests stage after the costs of fighting deflate. Thus, the same conflict that would generate one period of conflict of interests and two periods of armed contest in the null model, could result in two periods of conflict of interests and three periods of violence in the extended model. Again, in the case of multiple instances of power mediation, the difference between the two versions of the model would be greater.

Second, by contributing to other dyads through external support and power mediation, agents diverge from the settlements they achieved themselves. I assume that any agent may play the role of A or B with respect to another randomly assigned agent. In addition, this agent also plays the role of a third party with respect to all other dyads (except the one in which s/he is A or B). For instance, in a system of 200 agents paired up into 100 dyads, every single agent plays a role of either A or B with respect to one other agent and plays the role of a third party with respect to the remaining 99 dyads. Therefore, if agent number 76 pays the cost of reinforcement, the dyad in which agent 76 plays the role of A/B experiences a shift in capabilities toward agent 76's opponent. Depending on the previous distribution of benefits, such a shift is likely to lead to dissatisfaction with the previously acceptable status quo. In other words, acting as a third party generates shifts to capabilities that in turn "restart" previously settled conflicts. Consider another example: in a "world" of 200 agents paired up as 25 dyads, 50 agents play the role of either A or B, however every agent plays the role of a third party with respect to 50 dyads. In the former example, every instance of provision of support or power mediation results in a shift in capabilities in some other dyad, while in the latter example, third-party influence translates into a shift in capabilities only 1/4 of the time.

All three phenomena are the result of third-party influence on intrawar and postwar bargaining. Recall that external support operates through providing additional capabilities to the recipient which in turn changes the relative probability of victory which may either induce A to risk war by proposing a less generous offer or compel A to settle by giving up more of the disputed good than A would have otherwise. Thus, at face value, external support leads to more longer wars and more potential wars to be settled at the stage of conflict of interests. Since I assume that the resources provided through external support become a permanent part of the recipient's capability, reinforcement is able to solve an information problem when the shift in capabilities occurs toward imbalance. In contrast, when reinforcement moves relative power toward parity, it (re-) creates uncertainty and therefore prolongs conflicts. Thus, the system has a greater number of two-period wars in the extended model.[31]

In sum, conflicts restart for two reasons. First, the barriers to bargaining disappear, while a previously existing information problem remains unsolved. This is the effect of power mediation: it creates barriers to bargaining, which resumes once the barriers are removed. Second, the power dynamic changes in the dyad in such a way that another information

[31] Recall that in the simulation, a two period war is the longest possible war, because of the assumption that B may only be one of three types. Had I assumed that B may take on either of four types, the longest war would have been three time periods, etc.

problem arises when the sides renegotiate the new terms. A shift in power in itself is not necessarily damaging to peace, in fact a shift toward greater imbalance is supposed to be associated with a peaceful transfer of resources in favor of the growing side. However, when the shift moves the dyad toward greater parity, uncertainty may (re-)emerge.

The following discussion of the logic of how exactly power mediation and external support operate delivers two important theoretical implications that have emerged in the extended version of the model. While power mediation is in a sense "delaying the inevitable" by creating barriers to bargaining, it is possible for the renewed conflict of interests to avoid escalation: if external support shifts relative power in the dyad in such a way that uncertainty is resolved without resort to arms, then an imposed ceasefire created by power mediation is less likely to fail. However, if external support recreates uncertainty, an imposed ceasefire will be more likely to fail. The same dynamic is at work when one examines the effect of external support on war duration.

4.3.2 Scenarios that highlight the importance of third-party influence on dyadic bargaining

Within every dyad at peace, for every time period, the participants verify whether they are satisfied with the current status quo. If either side determines that the distribution of benefits does not match the relative power in the dyad, then a conflict of interests arises, which may or may not lead to an armed contest. I now describe two scenarios that may unfold in the extended model, under the condition that one of the sides is dissatisfied with the distribution of benefits in the dyad:

1. Power mediation inflates the costs of fighting;

2. Reinforcement either worsens or alleviates information problem.

Scenario 1: Power mediation inflates the costs of fighting. Consider a dyad of agents number 18 and number 67 paired such that 18 is assigned the role of side A and 67 is side B (hereafter, referred to as agents as A_{18} and B_{67}).[32] A_{18}'s capabilities are 0.02409 and B_{67}'s observable capabilities equal 0.02035,[33] and A's guesses about B's concealed capabilities include $\underline{b} = 0.0012; \hat{b} = 0.0035; \overline{b} = 0.0048$.[34] I have chosen a pair in which B's concealed capabilities are especially important as different beliefs of A will generate a different result with respect to parity and B's advantage. The belligerents' costs of fighting are $c_A = 0.007$ and $c_B = 0.005$.

Recall that the rest of the 148 agents in the system are considered third party agents with respect to this dyad. Each third party has capabilities, three randomly generated utilities to reflect his/her preferences for an instrument of influence on bargaining, and a preferred policy outcome. For the purpose of illuminating the logic of third-party power mediation on bargaining, consider an agent I_{145} who controls $m_{IT_{145}} = 0.30201$ and prefers $u_{IA_{18}} = .013$, $u_{IB_{76}} = 0.011$, and $u_{I_{\Pi}} = 0.971$. This third party may or may not act as a belligerent in some other dyad.

Given the randomly generated initialization parameter values for this dyad, A_{18} estimates that her probability of military victory is $\overline{p}_t = \dfrac{m_{A_{18}}}{m_{A_{18}} + m_{B_{67}} + \underline{b}} = \dfrac{.02409}{.02409 + 0.02035 + 0.0012} =$

[32] Refer to subsection 2.3.2 for a detailed explanation of the logic of the dyadic bargaining process when third parties cannot influence belligerents.

[33] The simulated capability data in the model approximate the empirical distribution of Composite Index of National Capability (CINC) scores as shown in Figure 2.5.Recall that agent capabilities in the model follow beta distribution with parameters $\alpha = 0.05$ and $\beta = 1.25$, such that the distribution of capabilities in the world approximates the empirical distribution of the CINC score (Singer, Bremer and Stuckey 1972).

[34] Recall that $\underline{b} < \hat{b} < \overline{b} < 1/3$ of m_B; in this case the program first generates a random value of $\overline{b}$ until the condition of $\overline{b} < 0.02035/3 = 0.0067833$ is met. Then the program generates random values for $\underline{b}$ and $\hat{b}$ until it finds values that are less than $\overline{b}$.

0.5278 if B is weak; $\hat{p}_t = 0.502$ if B is moderately strong; $\underline{p}_t = 0.4892$ if B is strong. For the purpose of this example let's assume that nature determines B's type as strong. Recall that side B is dissatisfied if his exogenously given proportion of the good $1 - \pi$ is less than what B could receive from fighting a war, which is $(1 - p_t) - c_B$, or $\pi > p_t + c_B$. A is dissatisfied with the status quo if she receives fewer benefits π than what A needs to give up to keep B satisfied, or $\pi < p_t + c_B$. Assuming that $\pi = 0.7$, B is dissatisfied, because he should control $(1 - \underline{p}_t - c_B = 1 - 0.4892 - 0.005) = 0.5058$ instead of 0.3 that he currently controls.

In the null model, the only way for A to appease B would be to give up a proportion of the disputed good that would match B's military strength, in other words, even if A correctly guesses B's type and gives up enough resources to appease even the strong type of B, a conflict of interests always involves re-division of the prize in question. In the extended model, A also first calculates her best guess offer to B, however after that, all third parties in relation to this dyad make decisions about reinforcement and power mediation, after which A recalculates her offer to B.

Assume that A's beliefs are $\phi_w = 0.2672$ and $\phi_m = 0.3929$. A's decision-making process is more involved as A calculates three values that describe A's "cutpoint" values for preferring $\underline{x}_t$ to $\hat{x}_t$, $\underline{x}_t$ to $\overline{x}_t$, and $\hat{x}_t$ to $\overline{x}_t$. A_{18} first calculates cutpoints 1 and 2 and then compares ϕ_w to the minimum of the two values.

$$\phi_w^{k_1} = \frac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B} - \phi_m = \frac{0.007 + 0.005}{0.502 - 0.4892 + 0.007 + 0.005} - 0.3929 = 0.090\ ,$$

$$\hat{\phi}_w^{k_2} = \frac{c_A + c_B}{\overline{p}_t - \underline{p}_t + c_A + c_B} - \phi_m \frac{\hat{p}_t - \underline{p}_t}{\overline{p}_t - \underline{p}_t + c_A + c_B} =$$

$$= \frac{0.007 + 0.005}{0.5278 - 0.4892 + 0.007 + 0.005} - 0.3929 \times \frac{0.502 - 0.4892}{0.5278 - 0.4892 + 0.007 + 0.005} = 0.1845.$$

$$\hat{\phi}_w^{k_3} = \phi_m \frac{c_A + c_B}{\overline{p}_t - \underline{p}_t} = 0.3929 \times \frac{0.007 + 0.05}{0.5278 - 0.4892} = 0.5801.$$

If A's belief that B is weak ϕ_w is less than or equal to the minimum of cutpoints 1 and 2, then A's best guess is that B is a strong type. A's $\phi_w = 0.2672$ is greater than the minimum of $\phi_w^{k_1} = 0.090$ and $\phi_w^{k_2} = 0.1845$, therefore B does not appear strong in A's estimation. If A's belief ϕ_w is greater than cutpoint 1 and is less than cutpoint 3, then A believes that B is moderately strong. A's $\phi_m = 0.3929$ is indeed greater than $\phi_w^{k_1} = 0.090$ and is less than $\phi_w^{k_3} = 0.5801$. Therefore, A believes that B is moderately strong. Finally, to demonstrate that A's beliefs are mutually exclusive, if A's belief ϕ_w is greater than the maximum of cutpoints 2 and 3, then A believes that B is weak. A's $\phi_w = 0.2672$ is less than the maximum of cutpoints 2 and 3, $\phi_w^{k_3} = 0.5801$.

Given this assessment, A first calculates that to appease a moderately strong B, she needs to offer $\hat{x}_t = \hat{p}_t + c_B = 0.502 + 0.005 = 0.507$. In the null model, given that a strong B would reject this offer, B's rejection would solve A's uncertainty, as only a strong type rejects $\hat{x}_t$ and after fighting one battle, A would propose $\underline{x}_t$, which would settle the dispute. In the extended model, however, third parties affect A's initial calculation by choosing to provide reinforcement or power mediation before A makes her first offer. In this example, I consider only three third parties, I_{39}, I_{91}, and I_{145}, out of the population of 148 outside actors. Third parties make their decisions in a random order.

Assume that in this run, I_{145} is chosen to make a decision first. Exogenously given utilities for certain bargaining outcomes determine whether I_{145} will calculate his/her expected utility of external support or of power mediation. Given that I_{145} clearly prefers immediate peace, $u_{I_\Pi} > u_{IA_{18}}$ and $u_{I_\Pi} > u_{IIB_{76}}$. Since I_{145}'s $u_{IA_{18}} = .013$ is greater than $u_{IB_{76}} = 0.011$, I_{145} calculates her $EU(PM) = \dfrac{u_{I_\Pi} - u_{I_A}}{|\pi' - \pi|} - \dfrac{\frac{m_A + m_B}{m_{IT}}}{m_A + m_B + b + m_{IT}} =$

$$\frac{0.971-.013}{.71-.7}-\frac{0.02409+0.02035+0.0035}{\frac{0.30201}{0.02409+0.02035+0.0035}}=\frac{0.958}{.1}-\frac{0.04794}{0.86}=9.58-0.055=9.52,$$

which is greater than 0, therefore the third party chooses to provide power mediation. It is no surprise that I_{145}'s expected utility from providing power mediation to this dyad is so high, as I_{145}'s preferred policy outcome is close to the current one, therefore the agent would want to preserve. Furthermore, the capabilities of this third party are much greater than the dyad's, therefore, this powerful actor may afford the costly action of mediation. As a result of I_{145}'s action, the costs of fighting for A and B become $c_A = 0.007 + \frac{0.30201}{0.02409+0.02035+0.0035+0.30201} = 0.007+0.86 = 0.867$ and $c_A = 0.005+0.86 = 0.865$.

Given these inflated costs, the formerly dissatified side in the dyad B_{67} reexamines whether he is still dissatisfied with the status quo of $\pi = 0.7$. B is dissatisfied, if $\pi > \underline{p}_t + c_B$ ($\underline{p}_t$ is used because B knows his type). Since $0.4892 + 0.865 = 1.3542$ is greater than .7, B is now satisfied. Similarly, since π of .7 exceeds her payoff in a potential military outcome even if she believes now that B is strong $\underline{p}_t - c_A = .4892 - .867 = -0.3778$, A is also satisfied with the status quo.

Third party pressure for peace that operates through inflating the costs of fighting prevents the sides not only from fighting. An important caveat to this mechanism is that such pressure or power mediation operates through making violence too expensive to bear, as opposed to solving uncertainty. If pressure is applied during the third loop of the simulation, after the sides have fought a battle, yet uncertainty still exists, the increased costs of fighting will make it increasingly likely that A proposes an appeasing offer of $\underline{x}_t$ even if without such pressure A would still risk war by offering $\hat{x}_t$. Thus, the mechanism of preventing war or making a warring dyad settle after the first battle is analogous in the model.

Inflation of costs may either precede armed conflict (prevention of onset) or follow the first battle (reduction of conflict duration) or follow the settlement (prevention of recurrence). I have discussed the prevention of conflict onset, now let a third party seek end of violence when the sides have fought a battle, and uncertainty still remains about whether B is a moderately strong or strong type (see step IV.4 of the detailed pseudocode). Recall that A compares her updated belief to the value that is derived from comparing A's utility of proposing the terms as if B is moderately strong vs. proposing the terms as if B is strong: $\phi_m^w > \dfrac{c_A + c_B}{\hat{p}_t - \underline{p}_t + c_A + c_B}$. If a third party inflates the costs of fighting, A will be increasingly likely to propose an appeasing offer of $\underline{x}_t$ which guarantees peace. As the new ϕ_m is recorded as ϕ_m^w, A will be able to reassess her offer once the costs of fighting decreases (given that the dyad will be at conflict of interests stage again). Therefore, when provided during an ongoing war, power mediation is likely to generate a ceasefire, which will have a higher chance of failing when the costs of fighting deflate, because power mediation does not manipulate A's beliefs.

If the sides have settled in the course of the simulation, a shift in dyad's relative capabilities may generate a conflict of interests to recur (see the scenario further below on this point). In empirical terms, this implies that the pressure that third parties use to prevent conflict onset should also be effective at preventing conflict recurrence: as long as the barriers to fighting are high, peace should be preserved. An empirical example of a continued pressure is the US's threats of withdrawing significant amounts of military aid if peace between Egypt and Israel is violated (Arena and Pechenkina N.d.). Another example of inflating costs of fighting would be peacekeeping with a mandate to engage or establishment of physical obstacles to fighting (e.g., demilitarized zones, separating walls). These third-party actions also do not solve the underlying conflict of interests, however make fighting costly enough that while these inhibiting factors are present, the sides choose not to fight.[35]

Importantly, unlike external support, power mediation prevents the sides from transferring resources (either during the stage of conflict of interests or during the armed conflict stage) to make the division of benefits more reflective of relative power in the dyad. Such pressure could be an effective instrument for keeping the sides from fighting, as long as a third party is motivated to keep the costs of armed conflict high. In the model, I assume that the costs of fighting stay inflated only for three time periods. Unless a shift in capabilities prevents the sides from going back to the conflict of interests stage, power mediation is followed by conflicts of interests and potential armed conflicts.[36]

Given that the costs of fighting need to be inflated by a third party, the motivation of outside actors is critical for the survival of imposed peace. The assumption of how long pressure may keep the costs of fighting prohibitively high is at the root of how one evaluates the effect of such pressure in a formal model. If one assumes that pressure for peace keeps the costs of fighting high for 10 periods, then such pressure would appear significantly more effective, than if the effect could last only 3 periods. However, such result is the direct outcome of a model assumption, rather than an emergent property of the model. Regardless of what one assumes such duration to be, the logic of this result is consistent with work

[35] Fortna (2004) demonstrates that peacekeepers and demilitarized zones have a pacifying effect; Fortna (2008) shows that peacekeeping is effecting at preventing civil war recurrence.

[36] While the aforementioned works explain that the same power mediator could seek to extend the barriers to fighting for a long time period, in the model, third-party decision making is done in a random order, implying that the same highly motivated third party is unlikely to precede some other outside actor in time period $t+4$. To reflect that the same interested third parties return to manage recurring conflict in the same dyad, I would need to change the order of third-party decision making in future iterations of the model. Although such a change could describe the empirical reality more closely, it would not uncover any new dynamic of third-party influence on bargaining, as the mechanisms of influence remain the same.

by Werner and Yuen (2005); Beardsley (2008, 2011). The authors argue and demonstrate empirically that imposed ceasefires fail more rapidly than those established without pressure, because imposed ceasefires are established prematurely before the true conflict is resolved. Once third-party pressure wanes, the sides resort to arms. While my model yields that this is a logical argument, an important challenge to this result emerges. Since I incorporate both support-providing and mediation-providing third parties in the same model, I derive that support-providing third parties actively change the stability of an imposed ceasefire, therefore the authors' argument is incomplete.

In this model, once the effect of pressure is over (3 periods), the dyad's participants reevaluate their satisfaction with status quo. Since power mediation does not encourage the sides to change the distribution of benefits, the previously existing conflict of interests reemerges. Once a dyad is at the stage of conflict of interests, both power mediation- and support-providing third parties make decisions about whether to pressure or reinforce. It is possible that another motivated provider of pressure for peace will keep the participants from fighting again, however, it is also likely that a third party motivated by a distributional outcome will consider investing in either side. If third parties reinforce the recipient by bringing the capabilities closer to parity, then they risk recreating uncertainty, thus motivating A to underestimate B's type. Alternatively, if third parties push the relative capabilities toward imbalance, then they eliminate uncertainty and therefore encourage peaceful transfers of resources and secure peace. This is a novel result as peaceful transfers of resources are an instrument through which external support establishes enduring peace (unless some external shocks to capabilities take place). Although it is probable that the same third party could first impose a ceasefire and then match the distribution of power and benefits (through reinforcement), in this model, third parties' preferences do not evolve and instead stay fixed during an entire simulation run. Therefore, the discussion further below about the binary nature of reinforcement (both pacifying and war-inducing) applies to post-settlement dyads as well, which means that settlements may either be undermined or strengthened by support-providing third parties.

Scenario 2: Reinforcement either worsens or alleviates the information problem.

Support is more common in the system as each third-party agent has two utilities that describe preferences for support and one that describes preference for power mediation. Consider a dyad A_2 and B_{111} with respective capabilities $m_{A_2} = 0.12$ and $m_{B_{111}} = 0.1$, including concealed $\underline{b} = 0.01, \hat{b} = 0.02, \overline{b} = 0.03$, and costs of fighting $c_A = 0.025$ and $c_B = 0.03$. A_2's beliefs are $\phi_w = .55$ and $\phi_w = .33$. In the first time period, the dyad starts at virtual parity: $\overline{p}_t = \dfrac{m_{A_{18}}}{m_{A_{18}} + m_{B_{67}} + \underline{b}} = \dfrac{0.12}{0.12 + 0.1 + 0.01} = 0.52173913043$; $\hat{p}_{=}0.50$; $\underline{p}_t = 0.48$. As a result, $\hat{\phi}_w^{k_1} = 0.7333 - .333 = 0.4003$; $\hat{\phi}_w^{k_2} = 0.57894736842 - .33 \times 0.210526 = 0.5094737$; $\hat{\phi}_w^{k_3} = .333 \times 1.375 = 0.45787$. Since A's belief $\phi_w = .55$ is greater than the maximum of cutpoints 2 and 3, which it is in this case, then A believes that B is weak, thus A will risk war by proposing $\overline{x}_t$. In the null model, depending on B_{111}'s type the conflict could potentially last two battles, if B_{111} happened to be strong.

In the extended model, however, third parties take turns to make decisions about reinforcement or power mediation. Consider I_{91}, who controls $m_{IT_{91}} = 0.40957$ and has the following preferences: $u_{IA_{18}} = .0129$, $u_{IB_{76}} = 0.951$, and $u_{I_{\Pi}} = 0.017$. Let I_{91}'s desired policy outcome be $\pi' = .15$; in other words, I_{91} seeks for B_{111} to control .85 of the good and A_2 to control .15. The optimal investment: $m_{I_{91B111}} = \dfrac{\pi' m_B + \pi' b - (1 - \pi')m_A}{(1 - \pi')} =$

0.56. The new m_{B111}=0.56+0.1=0.66, thus the new probabilities of A's victory are: $\overline{p}_t = \dfrac{m_{A_{18}}}{m_{A_{18}} + m_{B_{67}} + \underline{b}} = \dfrac{0.12}{0.12 + 0.66 + 0.01} = 0.1518987$; $\hat{p}_{=}0.15$; $\underline{p}_t = 0.14814814$. The updated probabilities result in new values of cutpoints for A's decision-making: $\hat{\phi}_w^{k_1} = 0.634$ $\hat{\phi}_w^{k_2} = 0.936171 - 0.01049674 = 0.92567426$ $\hat{\phi}_w^{k_3} = .333 * 14.66705 = 14.6670$. If A's belief that B is weak $\phi_w = .55$ is less than or equal to the minimum of cutpoints 1 and 2, then A's best guess is that B is a strong type. Therefore A will not risk war, and instead will give up $\underline{x}_t = \underline{p}_t + c_B = 0.14814814 + 0.03 = 0.178$.[37] Thus, shifting capabilities from parity toward imbalance makes war shorter, or as in this case prevents the violent portion of the conflict altogether.

This is an important takeaway from the simulation. Recall that imbalance in capabilities generates smaller distances in A's probability of defeating two different types of B. The greater the parity, the greater the gap between A's probability of defeating two different types of B. As this gap grows, the values of cutpoints 2 and 3 become smaller (as illustrated above) and, therefore, A becomes more likely to risk war by proposing $\overline{x}_t$ (or $\hat{x}_t$), because more randomly generated values of ϕ_m will pass the "threshold" of cutpoints 2 and 3. In contrast, the more closely the sides' capabilities approach preponderance of power, the smaller the difference in A's probability of defeating two different types of B. As this gap becomes smaller, the values of cutpoints 1 and 2 become larger (see the example above). The larger the values of cutpoints 1 and 2, the more likely A's randomly generated belief ϕ_w to fall under these cutpoints. If A's ϕ_w is less than or equal to the minimum of cutpoints 1 and 2, then A believes that B is strong, and therefore will appease strong B by proposing $\underline{x}_t$,

[37] Again, a third party has achieved even a better outcome than she was seeking, because the optimal investment calculation does not take into account the cost of fighting. This shortcoming will be removed in future versions of the model.

which B always accepts. Therefore preponderance of power in my model is associated with greater probability of peace and parity is associated with greater probability of war. Third parties may manipulate the distribution of capabilities in the dyad and therefore precipitate war or ensure peace. In the scenario above, I_{91} has solved the information problem in the dyad by shifting the distribution of capabilities so far toward preponderance that A could transfer the necessary resources without risking war.

Given that the randomly generated values of ϕ_w and ϕ_m are critical in both the values of cutpoints and whether A's assessment of B's type will change after reinforcement, a question arises: if A's beliefs happen to be stacked in such a way that regardless of reinforcement A will keep her original offer, does reinforcement still make sense for a third party? The answer is yes. A third party is still going to shift the policy outcome closer to her/his desired outcome, therefore investment translates into sought policy regardless of whether the battle dynamics change. Consider the aforementioned dyad of A_{18} and B_{67}. Let $\pi = 0.5$. Let B know from nature that he is strong, and therefore is dissatisfied, as he should control $(1 - (\underline{p_t} + c_B)) = (1 - 0.4942) = 0.5058 \approx .51$ of the prize, as opposed to 0.5. In other words, B seeks $\pi = .49$. Let A's beliefs be $\phi_w = 0.05$ and $\phi_m = 0.1$ for the same dyad A_{18} and B_{67}. While the sides are virtually at parity the beliefs were randomly generated such that it is extremely likely for A to believe that B is strong no matter what. These beliefs, together with the probabilities of victory and costs of fighting, generate $\hat{\phi}_w^{k_1} = 0.384$, $\hat{\phi}_w^{k_2} = 0.212$, and $\hat{\phi}_w^{k_3} = 0.031$. If A's belief that B is weak ϕ_w is less than or equal to the minimum of cutpoints 1 and 2, then A's best guess is that B is a strong type. In this scenario, $\phi_w = 0.05$ is less than both $\hat{\phi}_w^{k_2}$, therefore A estimates that B is strong. In the null model, if A decides to propose the most generous offer of $\underline{x}_t = \underline{p}_t + c_B = 0.4892 + 0.005 = 0.4942$, B

always accepts and the conflict of interests gets resolved through the reallocation of resources proportionally to the participants' relative power. In the extended model, however, third parties' influence may change both the dynamic and the outcome of bargaining. In this scenario, a support-providing third party affects only the outcome of bargaining, while the dynamic is preserved.

Assume that in this execution of the model, I_{39} is chosen to make his decision first. Given that I_{39} prefers A's victory ($u_{IA_{18}} > u_{IB_{76}}$ AND $u_{IA_{18}} > u_{I_{\Pi}}$), I_{39} is likely to be considering external support to A. Besides exogenously given utilities for certain bargaining outcomes, I_{39} also needs to determine whether he possesses the necessary resources to consider his expected utility of external support. Recall that m_{I_A} needs to be within 30% of I_{39}'s total capabilities. Let I_{39}'s preferred policy outcome in this example be $\pi' = .51$, which is greater than what A currently controls. Since I assume that a potential supporter is motivated to maximize his preferred policy outcome through increasin the recipient's probability of victory to the "necessary" level, $m_{I_{39A18}} = \dfrac{\pi' m_B + \pi' b - (1-\pi')m_A}{(1-\pi')} =$ $\dfrac{0.51 m_{B_{76}} + 0.51\bar{b} - 0.49 m_{A_{18}}}{0.49} = \dfrac{0.51 \times 0.02035 + 0.51 \times 0.0048 - 0.49 \times 0.02409}{0.49} = 0.0021,$ which is well within $0.3 \times m_{IT_{39}} = 0.3 \times 0.34731 = 0.104193$. Therefore, I_{39} calculates expected utility from reinforcing side A: $EU_{I_{39}}(ES_{A_{18}}) = (u_{I_A} - u_{I_B})(\pi' - \pi) - m_{I_A} =$ $(.967 - 0.015)(.51 - .5) - 0.0021 = 0.0931$, which is greater than 0, therefore I_{39} contributes $m_{I_{iA}}$ of 0.0021 to A_{18}. This means that, A_{18}'s capabilities change to $m_A = 0.02409 + 0.0021 =$ 0.02619, which, in turn, affects A's relative probability of victory: $\overline{p}_t = \dfrac{m_{A_{18}}}{m_{A_{18}} + m_{B_{67}} + \underline{b}} =$ $\dfrac{0.02619}{0.02619 + 0.02035 + 0.0012} = 0.54859$; $\hat{p}_{=}0.52338$; $\underline{p}_t = 0.510128$.

After receiving reinforcement, A_{18} also recalculates the cutpoint values, which now are: $\hat{\phi}_w^{k_1} = 0.375$, $\hat{\phi}_w^{k_2} = 0.457$, and $\hat{\phi}_w^{k_3} = 0.020$. Even after the reinforcement A_{18} still believes that B_{76} is strong, therefore A_{18} will propose an appeasing most generous offer of $\underline{x}_t$, however, after reinforcement, the size of $\underline{x}_t$ increases, therefore I_{39}'s external support allows A_{18} to keep more of the prize. $\underline{x}_t = 0.510128 + 0.005 = 0.51512$, which is greater than the previously calculated most generous offer of 0.4942 (without reinforcement). Recall that I_{39}'s preferred policy outcome is $\pi' = 0.51$.[38] I_{39}'s reinforcement is able to improve the policy outcome such that the division of benefits becomes more consistent with the third party's preference, even though the dynamic of bargaining (i.e., the conflict of interests is still resolved without violence) is unchanged.

It is important to note that investments on both sides of the dyad are common (determined randomly, given the exogenously given utilities) and undermine the effect of reinforcement on the opposing side. Thus, a portion of each I_i's investment that affects uncertainty in the dyad is "net" reinforcement.

Finally, the supply of external support affects both further provisions of external support and power mediation by other third parties. Recall that all actors make decisions in random order. Once I_i chooses to reinforce/mediate, capability/cost terms of the dyad are updated immediately, such that the next randomly selected I_{i+n} is able to view those changes and take them into account. For instance, after I_i invests in A, I_{i+n}'s reinforcement of B will be that much more expensive (assuming that preferred policies by I_i and I_{i+n} diverge), and therefore marginally less likely. Similarly, if I_i reinforces either side of the

[38] The final settlement slightly exceeds I_{39}'s target division of benefits, because I assume that third parties do not take into account the costs of fighting when calculating the necessary investment. This shortcoming will be revised in the future version of the model.

dyad first, the recipient updates her/his capability, and then I_{i+n} provides power mediation, I_i makes power mediation less attractive for I_{i+n}, because the greater amount of dyad's total capability increases the cost of power mediation for I_{i+n}.[39] If, however, I_{i+n} provides power mediation to the dyad, both sides' costs of fighting increase, and then I_i chooses to reinforce either side, power mediation will not prevent I_i from reinforcement, however its impact will be delayed for as long as the costs of fighting stay inflated. Given that power mediation affects the costs of fighting, third parties that are motivated by reinforcement, are unable to consider those costs in the current version of the extended model.[40] This, in turn, implies that as long as the sides are satisfied with the status quo, which is the most common outcome of power mediation, the relative shift in capabilities does not affect the bargaining dynamic, as none takes place. Once the barriers to fighting wane, and the dyad finds itself at a conflict of interests again, the net shift in power affects A's estimation of B's strength and thus the size of A's offer. In other words, power mediation delays reinforcement's effect on the likelihood of armed conflict.

4.3.3 Analysis of the extended model with *SimGA*

The software application *SimGA* introduced in chapter 3, provides communication between a simulation model and a genetic optimization algorithm, such that the genetic algorithm evolves the initialization parameter values of the model until the performance

[39] Recall that the cost of power mediation involves $\dfrac{\dfrac{m_A + m_B}{m_{IT}}}{m_A + m_B + b + m_{IT}}$.

[40] In the next model iteration, the optimal size of investment needs to be estimated as $m_{I_A} = \dfrac{\pi' m_B + \pi' b - c_B - (1 - \pi') m_A - c_A}{1 - \pi'}$, which will effectively deter reinforcement as long as the barriers to fighting are present.

of the model achieves certain criteria of fitness. In chapter 3, I verified the null version of the model by letting the genetic algorithm evolve the initialization parameter values to maximize the amount of violent conflict in the system. Recall that in the null model, capabilities do not change over time, therefore, initial distribution of capabilities in the dyad is critical: the closer the sides are to parity, the greater the difference between A's probability of defeating two different types of B and the greater the chance that A will choose to risk war.[41] Similarly, the greater the initial values of costs of fighting, the less likely the sides are to fight. In this section, I employ the genetic automated search to understand what initialization parameter values are associated with greater amounts of armed conflict in the system and greater amounts of peace in the system. The fundamental difference between the two versions of the model is that -unlike in the null model- capabilities change throughout the existence of the extended version of the "world." This leads to no relationship between the initial distribution of capabilities in the dyads and amount of conflict or peace a system will generate. The major drivers of amount of conflict and peace in the extended model are the preferences of third parties to mediate or support.

Just like the null model, the extended version may be executed in a batch mode with initialization parameter values set via a parameter file and numerical output recorded in the form of one line summaries per model execution. To create a partial parameter space, the genetic search program uses a genetic algorithm to evolve new versions of initialization parameter values, creates and then feeds parameter files into the simulation model, runs the simulation model, records the values of the emergent characteristics of interest, and, finally,

[41] It is important to note that sides that are almost at parity are more likely to fight than the sides that are at exact parity, because of the concealed capabilities that B is assumed to possess in the model.

reads in the information from numerical output summaries generated by the simulation model to evaluate and evolve the following generation of candidate solutions.[42]

In contrast, to examining the evolution of conflictual and peaceful interactions over simulation time as described in section 4.3, here, I explore only the relationship between initialization parameter values and the values of the outcome variable: count of armed conflicts in the system. In other words, I want the genetic algorithm to yield what combinations of initialization parameter values can produce the highest possible number of battles and what combinations lead to the most number of time periods, when no dyad experiences a conflict of interests or an armed conflict. Given that I am looking at systems with 200 agents and 50 dyads that are "cut off" after 70 time periods,[43] the highest possible number of conflicts would result from each dyad having an armed contest every time period. Given that each war ends in two battles (because of the assumption that B may only be either of the three types), for a war to recur, there has to be at least one round of conflict of interests, therefore every third round should involve a conflict of interests, as opposed to an armed conflict: $(50 \times 70)/3 = 1166.7$ is the highest possible number of armed conflicts in the system. This is the target value, against which the candidate solutions are evaluated.

Recall that *SimGA* program treats the numerical one-line summary of a simulation run as one function evaluation. The analysis of the null model with *SimGA* requires over 1.6 million function evaluations, which is a great reduction in computational time and power compared to over 179 million function evaluations (i.e., simulation runs) required to create

[42] See a detailed outline of the genetic automated search program in Figure3.3.

[43] I have done executions of both the null model and the extended model that cut off the model at 400 time periods and the conclusions described in sections 2.3.2 (stable global peace) and 4.3 (conflicts recur throughout the life of the system) hold.

an exhaustive parameter space of the null model. In the extended model, after the inclusion of four additional initialization parameters each with 10 discrete values, an exhaustive space would include over 179 billion function evaluations. To avoid such computational commitment, I employed automated genetic search to create a partial parameter space.

Table 4.4 lays out a potential parameter file that one may want to use to generate an exhaustive parameter space of the extended model. The parameters unique to the extended model include characteristics of third parties: capability and utility values that describe third party preferences for bargaining outcomes (see Table 4.2 for a detailed description). The illustrative runs of the model discussed in section 4.3 record numerical output as time series cross sections of dyadic interactions. Such detailed numerical output is important for understanding major model dynamics. In addition, detailed output data allow the analyst to verify if the recorded interactions that emerge from the computer model are consistent with model assumptions. However, such detailed data are difficult to summarize across multiple sets of model output. Therefore, my analysis of the model with genetic automated search is manageable when the genetic search program relies on the batch mode-style numerical output from the simulation model, i.e., one line summary per one model execution. This necessity explains my decision to model dyads that have the same capability distributions, costs of fighting, and beliefs that they hold. The differences among the dyads in the batch mode of the null model include the type of B, the distributions of benefits, and other randomization steps throughout the algorithm of the model. In the extended model, these differences also include the capabilities of third parties and third parties' preferred policy outcomes for each dyad. To satisfy the requirement that 100 third parties do not participate in the dyadic interactions (which is the case in a system with 200 agents and 50 dyads), third party capability is included as an initialization parameter, along with third party utilities from dyadic bargaining outcomes. These differences between the two versions of the model are evident in the structure of candidate solutions for each version. A candidate solution in this context is a string of initialization parameter values. Table 4.3 compares the structure of candidate solutions in the null model and the extended model: the extended version of the model adds four new initialization parameter values that describe the added behavior in the system: third party influence.[44]

As outlined in section 3.3.1, the genetic search program starts each run with randomly generated candidate solutions and evaluates them according to their fitness. Every candidate solution (i.e., the values of the initialization parameters) are recorded to file together with the program's evaluation of that solution's fitness. The program ends its run and begins another one with a new random population when the sought target is achieved. In some cases, the target is not achieved in 1,000 generations, which is the default threshold for ceasing the program's execution (denoted as t_{max} in the discussion of *SimGA* pseudocode). The goal of imposing a cut off threshold is to avoid wasting computational time when the algorithm is not improving.[45]

I have carried out two attempts to maximize the count of armed conflicts in the extended model. In the first attempt, I omitted the third-party characteristics from the candidate solutions, making candidate solutions almost equivalent to those used in the null model (see column 1 of Table 4.3). The only difference being the added m_{IT} parameter to

[44] Also, see Tables 3.2 for an actual example of a candidate solution.

[45] The value needs to be scaled with the difficulty of the problem. In the null model, the longest number of generations required to locate the maximum was 132, assuming that the extended model is almost tenfold more difficult to solve, I use a threshold of 1,000 generations.

specify the initial capability of third parties. This omission allowed me to model a system with diverse third party preferences (i.e., the values of u_{I_A}, u_{I_B}, u_{I_Π} were different for every member of the system). When third parties have divergent preferences that cancel out possible shifts toward more conflict or peace in the system, the *SimGA* program aborted its execution after 1,000 generations in each of the 500 runs without having located the sought target of 1166.7. Given that each generation has 50 solutions and each solutions requires 50 executions of the agent-based model, a total of 2.5 million function evaluations were recorded for this exercise (see Table 3.3 for more details about the default settings of the genetic search application).

Panel (a) in Figure 4.4 shows the data from a single run of the genetic search application (out of 500 runs executed). Panel (a) charts the time series of the best fitness scores over the course of 1,000 generations. Insignificant improvements are followed by insignificant deteriorations, and the overall trend is that when third-party preferences are diverse, the simulated international system cannot achieve greater than the average of 539.1521 armed conflicts in 70 time periods.[46] The average outcome of armed conflicts in the system with 50 dyads and diverse third party preferences is 205.5122. Again, the reason why so many conflicts can be achieved is because imposed ceasefires may fail and thus generate recurring conflict, in addition to constant capability transfers among agents in the system as they reinforce and pressure other agents to settle. The total amount of capability in the system decreases over time due to power mediation being costly and the resources paid not being

[46] Recall that each candidate solution is assessed by executing the agent-based model 50 times and then a mean of model outcomes is evaluated as that solution's value of the outcome variable.

transferred to other agents. Another source of fluctuation of capability of each agent is receiving or providing support. Given that third party behavior constantly upsets status quo in the dyads, conflicts restart.

In my second attempt to maximize the amount of armed conflict in the system, I used the candidate solutions structured as shown in column 2 of Table 4.4, such that the genetic automated search could evolve the preferences of third parties. It is important to note that in this version of the model, all third parties share the same values of utilities for bargaining outcomes with respect to every dyad. For instance, a high value of u_{I_A} implies that every third party has a high preference for reinforcing side A in every dyad. Again, after 500 runs, the genetic search program could not achieve the sought 1166.7 armed conflicts, the maximum possible in this system. However, the results are improved compared to the initial attempts. Panel (b) of Figure 4.4 shows a single run of the genetic search program. Since the program has not located the target, it ran for 1,000 generations before cutting off the process of evolving the initialization parameter values. Panel (b) shows the time series of best fitness scores for each of 1,000 generations. Unlike in panel (a), there is modest improvement in the fitness score for the initial 56 generations of the run. Due to the small scale, the progress in best fitness scores from one generation to the next appears minuscule (this is due to nonlinearity in the formula of fitness score, and much greater values of dependent variable compared to the null model). However, the time series in panel (b) demonstrates that letting the genetic algorithm evolve third-party utilities for bargaining outcomes changes the sought fitness of the system: the best result across the 500 runs of the genetic search program is 972.9015 armed conflicts. However, the mean number of conflicts almost has not improved dramatically: 331.3712.

To understand what initialization parameter values are associated with higher volumes of conflict in the system, I have constructed the same variables examined in chapter 3 to describe the difference in A's likelihood to win militarily against B: $Difference_{\overline{p}_{t=1}, \hat{p}_{t=1}} = \overline{p}_{t=1} - \hat{p}_{t=1}$. $Difference_{\hat{p}_{t=1}, \underline{p}_{t=1}} = \hat{p}_{t=1} - \underline{p}_{t=1}$. $Difference_{\overline{p}_{t=1}, \underline{p}_{t=1}} = \overline{p}_{t=1} - \underline{p}_{t=1}$. Recall that the simulation model implies that greater difference in A's probability to defeat two different types of B is associated with A being more likely to risk war. As these variables are based on the initialization parameter values, they reflect the said differences in the first time period of the simulation. In subsection 3.3.3, I have demonstrated that in fact this relationship is evident in the partial parameter space created with the genetic search program coupled with the null model. Table 4.5 presents Pearson correlation coefficients for differences in A's probability of defeating two different types of B at time 1 and measures of candidate solutions' performance for the partial parameter space created with *SimGA* coupled with the extended model. Unlike in the analysis of the null model, there is no statistically significant relationship between the difference in A's probability of defeating different types of B in the first period of the simulation and the amount of armed conflict the system will reach. Furthermore, the size of the correlation coefficients is deflated as well compared to the null model's analysis shown in Table 3.4. This is not surprising: in the extended model, the capabilities and costs of fighting shift constantly due to third-party influence and, therefore, A's probability of defeating B changes over time. This result underscores the fundamental difference between the null and the extended versions of the model.

I have also examined the relationships between the preferences of third parties and measures of candidate solutions' performance for the partial parameter space created with *SimGA* coupled with the extended model. Recall that these preferences do not change over time, therefore these variables are more informative about the dynamics of the interactions than differences in A's probability of victory at the initial time period of the model. Table 4.5 shows a highly significant negative correlation level (in terms of both substance and statistical significance) between third-party preference for immediate end of violence u_{I_Π} and amount of armed conflict in the system. This is not surprising as high preference for peace leads to power mediation, which is effective at imposing a ceasefire. In contrast, there are positive correlations between utilities from supporting A or B and armed conflict. The size of the coefficients is not large for these relationships. A substantively weak relationship is consistent with the discussed logic of external support: it may shift the capabilities such that the information problem is either alleviated or exacerbated. While a substantively negligent effect would have been more consistent with the logic discussed in section 4.3, examination of correlations with the volume of peaceful interactions and the amount of conflict of interests will convey a fuller picture of the effect that external support has on system level outcomes.

4.4 Conclusion

In this chapter, I have developed and analyzed an extension to the agent-based model introduced in chapter 2. The fundamental difference between the two versions of the model is that the extended model includes third-party influence on bargaining interactions during the phase of conflict of interests and during the armed conflict phase. I assume that third party behavior may be grouped as power mediation and external support. Power mediation operates through inflating the costs of fighting to force a settlement, while external support shifts capabilities in the dyad such that a third party's preferred policy outcomes may be achieved more easily. I derive that the effect of power mediation is consistent with the current literature: pressure for peace delays the inevitable because it does not address the underlying information problem in the warring dyad. However, in my model, those dyads that have returned from an imposed ceasefire as dissatisfied with the status quo (once the effect of power mediation ends) may avoid violence if external support shifts the balance in the dyad further toward preponderance, which in turn makes it easier for the offer-making side to propose appeasing terms or realize her own predominance. Thus, we would expect peaceful transfers of resources in those cases. This pacifying effect of external support has been overlooked as scholars have mainly focused on whether the recipient is the target or initiator (in interstate conflict) or rebels vs. government (in civil wars). The effect of external support is not always pacifying, however. When a third party's policy preference motivates him to shift the balance of capabilities toward parity in the conflictual dyad, then the difference between A's probability of defeating two different types of B grows and A is more likely to risk war. In sum, the effect of external support on war duration is bimodal: when eliminating the information problem, support is pacifying, while it makes wars longer if support exacerbates uncertainty in the dyad. Furthermore, if power mediation's imposed barriers to fighting are temporary then a dyad is likely to go back to the conflict stage once the barriers disappear. However, in some cases repeated instances of power mediation prolong its effect significantly, and in other cases, supporters may eliminate the information problem and therefore encourage peaceful transfers of resources without military contest.

Fig. 4.1. Counts of peaceful dyads, dyads that experience a conflict of interests, and dyads at armed conflict shown over time in a single simulation run

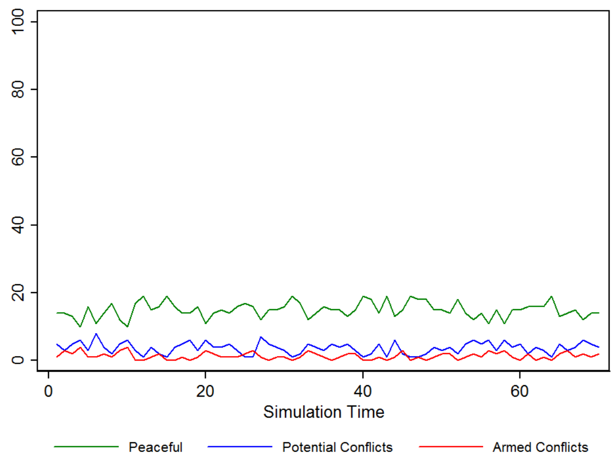

(a) "World" system of 20 dyads

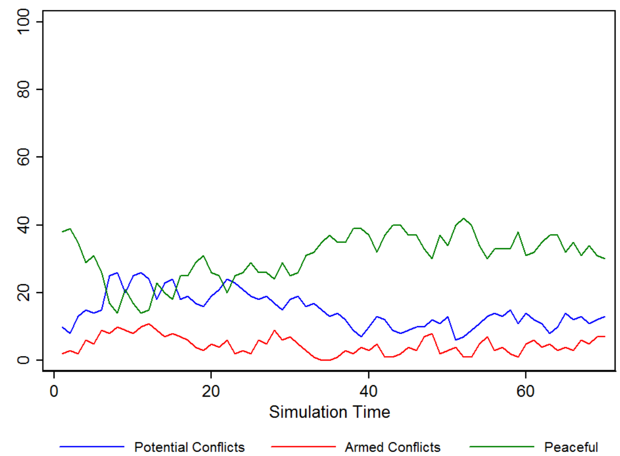

(b) "World" system of 50 dyads

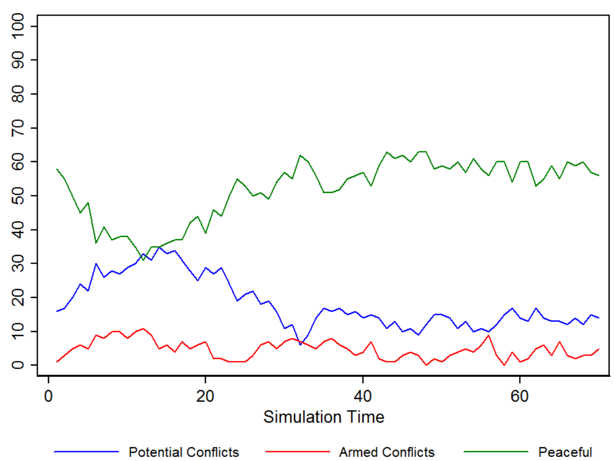

(c) "World" system of 75 dyads

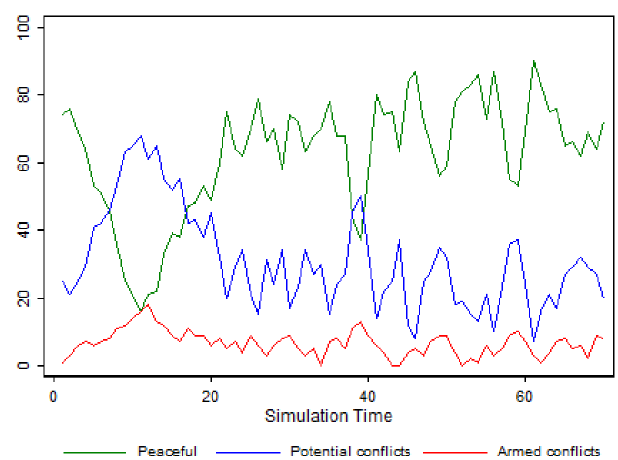

(d) "World" system of 100 dyads

Note: regardless of the number of dyads in the artificial international system the major trend remains that the volume of conflict is much greater in the extended model than in the null model. More specifically, conflicts of interests (and thus armed conflicts) keep recurring in the extended model, because imposed peace settlements are sometimes unstable and providing resources to other dyads undermines third parties' settlements. For comparison, consider the time series for the null model in Figure 2.1.

Fig. 4.2. Absolute changes in A's beliefs in 50 dyads over time in a single simulation run of the extended model

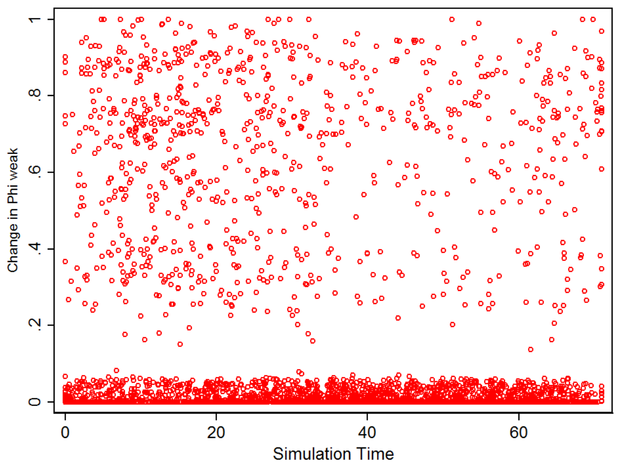

(a) Change in A's belief that B is weak

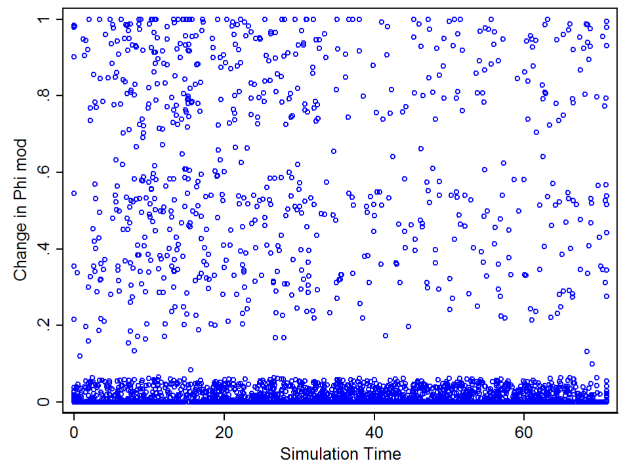

(b) Change in A's belief that B is moderately strong

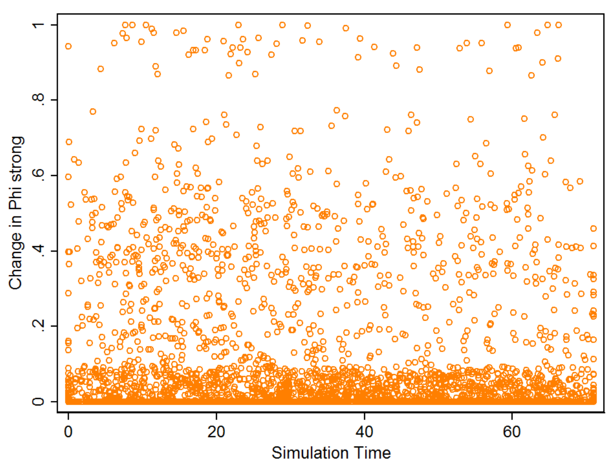

(c) Change in A's belief that B is moderately strong

Note: the graphs chart absolute differences between A's beliefs about B's strength in time period t from A's beliefs in time period $t-1$, i.e., $\phi_t - \phi_{t-1}$ in the extended model. For comparison, consider absolute changes in A's beliefs in the null model in Figure 2.3.

Fig. 4.3. Absolute changes in distribution of benefits, π, in 50 dyads over time in a single simulation run of the extended model

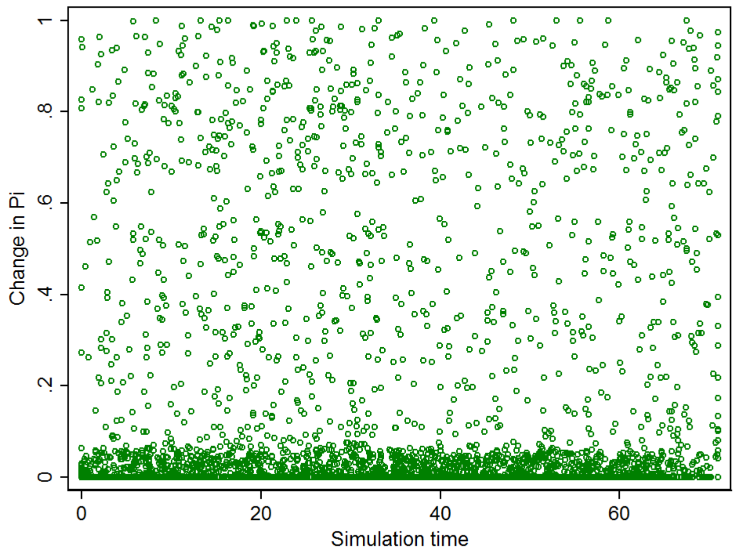

Note: The circles represent absolute changes in the division of the disputed good in each dyad between t and $t-1$, i.e., $\pi_t - \pi_{t-1}$. For comparison of how π changes in the null model, see Figure 2.4.

Fig. 4.4. Output from *SimGA* when coupled with the extended model. Maximizing armed conflict in the system

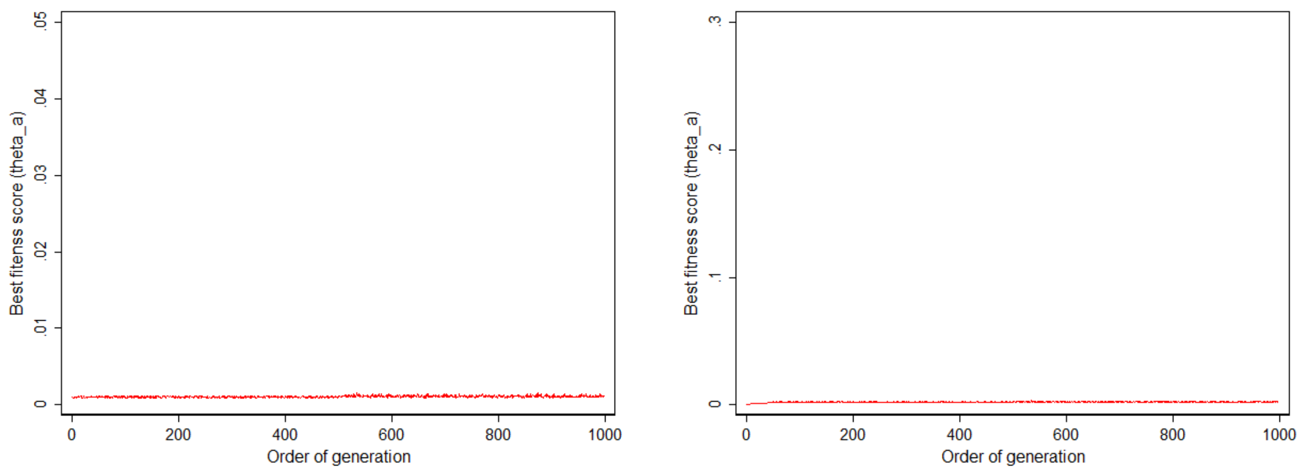

(a) Best fitness scores θ_a per generation over one run of SimGA (cuts off at 1,000 if target not found) – third party preferences are divergent

(b) Best fitness scores θ_a per generation over one run of SimGA (cuts off at 1,000 if target not found) - third party preferences are coordinated

Note: The graphs demonstrate that when third party preferences are divergent (panel A), the system will not generate as many armed conflicts as it will third parties having unified preferences (panel B). The size of the fitness scores is higher in panel B, even though none of the setups matches the sought fitness of over 1600 armed contests in 70 time periods.

Table 4.1. Additional initialization parameters in the extended model

Parameter	Code Reference	Range	Default Value	Interpretation
π'	`pi I`	$[0, 1]$	Uniformly distributed, SD=0.288	I_i's preferred division of benefits for a given dyad
m_{IT}	`caps I*`	$[0, 1]$	Beta distributed, $\alpha = 0.05$, $\beta = 1.25$	I_i's total capabilities
u_{I_A}	`utility A vic`	$[0, 1]$	Uniformly distributed, SD=0.288	I_i's utility from A's victory
u_{I_B}	`utility B vic`	$[0, 1]$	Uniformly distributed, SD=0.288	I_i's utility from B's victory
u_{I_Π}	`utility Peace`	$[0, 1]$	Uniformly distributed, SD=0.288	I_i's utility from immediate end of violence

Note: See Table 2.1 for a list of initialization parameter values that are present in both the null and the extended versions of the model.

Table 4.2. Additional characteristics and quantities of interest of actors introduced in the extended model

Additional characteristics of actors	
m_{IT}	I_i's total capabilities.
m_{I_A}	The amount of I_i's capability that I_i considers to invest in A, to increase the probability of A's victory. $m_{I_A} \in (0, 0.3m_{IT})$.
m_{I_B}	The amount of I_i's capability that I_i considers to invest in B, to increase the probability of B's victory. $m_{I_A} \in (0, 0.3m_{IT})$.
m_{IPM}	The cost to I_i of providing power mediation.$m_{IPM} \in (0, 0.3m_{IT})$.
u_{I_A}	I_i's utility from A's victory; exogenously given.
u_{I_B}	I_i's utility from B's victory; exogenously given.
u_{I_Π}	I_i's utility from immediate settlement; exogenously given.
π'	I_i's preferred division of the disputed good in a given dyad; exogenously given.
$EU_{I_i}(ES)$	I_i's expected utility from providing external support to either side of a given dyad.
$EU_{I_i}(PM)$	I_i's expected utility from providing power mediation to a given dyad.

Notes: See Table 2.2 for a list of characteristics present in both the null and the extended versions of the model.

Table 4.3. Comparison of candidate solutions for *SimGA* in the null and extended versions of the model

1 Null version	2 Extended version
m_A at t=1	m_A at t=1
m_B at t=1	m_B at t=1
ϕ_w at t=1	ϕ_w at t=1
ϕ_m at t=1	ϕ_m at t=1
$\underline{b}$	$\underline{b}$
$\hat{b}$	$\hat{b}$
$\bar{b}$	$\bar{b}$
c_A	c_A
c_B	c_B
	m_{IT} at t=1
	u_{I_A}
	u_{I_B}
	u_{I_Π}

Note: Some of the parameters are labeled as a certain value at t=1, because their values change in the course of the simulation, while other parameters do not have time specifications, as they stay the same.
All values are doubles with six digits after the decimal point, within (0, 1). An example candidate solution is shown in Table 3.2. Cost and concealed capability parameters have threshold limitations as described in Table 2.3.

Table 4.4. Sample parameter file to execute exhaustive search of the extended model

Parameter	Range	Discrete values used	Number of combinations
		Total of unique runs in the null model =	17,968,500
m_{IT}	$(0, 0.4]$	0.0001, 0.001, 0.05, 0.1, 0.15, 0.2, 0.25. 0.3, 0.35, 0.4	10
u_{I_A}	(0, .99]	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.99	10
u_{I_B}	(0, .99]	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.99	10
u_{I_Π}	(0, .99]	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.99	10
		Total of unique runs in the extended model=	179,685,000,000

Note: See Table 2.4 for discrete values of the initialization parameters shared between the null and the extended model.

Table 4.5. Pearson correlation coefficients for certain initialization parameter values and measures of candidate solutions' performance in the extended model

	θ_a	p-value	T_a	p-value
$Difference_{\overline{p},\hat{p}}$ $at\ t = 1$	0.13	0.2034	0.29	0.140
$Difference_{\hat{p},\underline{p}}$ $at\ t = 1$	0.18	0.783	0.09	0.342
$Difference_{\overline{p},\underline{p}}$ $at\ t = 1$	0.23	0.160	0.12	0.103
u_{I_A}	0.45	0.000	0.42	0.000
u_{I_B}	0.39	0.000	0.35	0.000
u_{I_Π}	- 1.89	0.000	- 1.75	0.000

Note: N=2.5 million. The table shows insignificant correlation between the differences in A's probability of defeating two different types of B in the first time period of the simulation and the candidate solution generating a high value of the simulation response variable (T_a) and thus being assigned a high fitness score (θ_a). This result is due to capability and costs of fighting changing in the course of the duration of extended model. In contrast, third party preferences are highly statistically significant: preference for immediate peace is negatively related to the volume of conflict in the system. Preferences for reinforcement have weak positive relationships with conflict level.

Chapter 5

Empirical Analysis of the Extended Model's Implications

I use observational data to test empirically two of the extended model's implications developed in chapter 4. I find that civil war duration after 1945 is affected by third party reinforcement consistent with the extended model: net reinforcements toward parity prolong war, while net reinforcements in the direction of imbalance shorten it. Using a sample of ceasefires established after interstate wars in the twentieth century, I also demonstrate that third-party reinforcements shape peace duration in accordance with the extended model. Shifts toward imbalance prolong peace, while net reinforcements that increase parity shorten it.

5.1 Hypotheses Developed in the Extended Model

External support may shape conflict dynamics through manipulation of uncertainty within the conflictual dyad. By shifting a balance of capabilities within a warring dyad, third parties may either lessen or aggravate the information problem that is at the root of fighting. When third parties shift the balance of power marginally toward preponderance, they make it easier for the offer-making belligerent to guess the type of her opponent. In contrast, when the balance of capabilities is shifted marginally toward parity, the advantage of defeating a weaker opponent as opposed to a stronger opponent appears greater and, therefore, the offer-making belligerent is more willing to risk another battle, and thus war is more likely to continue. The result that parity is associated with war onset is well established in international relations (Reed 2003). My model generates the following contributions with respect to war duration:

- Pursuing their desired policy outcomes, third parties sometimes shift relative power in the direction of parity, thus, creating greater uncertainty and encouraging more fighting. Other times, outside actors create greater imbalance, thus, reducing uncertainty and shortening war;[1]

- The examination of my model's general scenarios of multi-party interaction yields that it is shifts in the power structure of a conflictual dyad that shape war duration as opposed to identities of the recipients of support.[2] In the empirical literature on third party support in civil war, the major question has been whether providing reinforcements to government or opposition lead to peace. Scholars of duration of international conflict have focused on the reinforcement of either target or initiator. I ignore the identity of the recipient and demonstrate that focusing on how reinforcement affects the power structure in the dyad yields a more general empirical model of third-party influence on conflict duration.

To emphasize the importance of power structure as opposed to the identity of belligerents, I test across samples of inter- and intra-state wars, whether:

[1] While it has been established that preponderance of power leads to shorter wars (Bennett and Stam 1996), I focus on whether third-party reinforcements that lead to preponderance shorten war.

[2] The systematic analysis via the genetic algorithm software supports the formulated hypotheses, however for the reasons enumerated in Chapter 4 the current version of the SimGA software is not able to be a solo instrument for analysis of relatively complex computational models. Once a more sophisticated version of the text interpreter is added to the SimGA code, the software would be able to interpret time-series cross-sectional numerical output, which is critical for the comprehensive analysis of the extended model in this project.

Hypothesis 1: Net reinforcements toward greater imbalance shorten war and net reinforcements toward greater parity prolong war.

In addition, my model yields that the effect of external support is muted by power mediation.[3] If one assumes that preferences for policy outcomes do not change rapidly over time, then external supporters remain motivated to reinvest when the effect of mediation wanes (or, the effect of support materializes with a delay as it is shown in the extended model). Whether power mediation lessens th effect of external support or prevents support from being provided altogether. I do not test for whether power mediation discourages external support in conflicts or whether the presence of power mediation undermines the effect of external support on war duration; this remains an important question for my future work.

As conflicts of interests reemerge in the extended model, third parties decide (again) whether they will provide external support or power mediation to the conflictual dyad. Just as net reinforcement toward parity may encourage the sides to risk another battle during war (i.e., worsen the information problem), such net reinforcement may also provoke the sides to reengage militarily even after the ceasefire is signed, as the benefits of defeating two different types of B appear that much more attractive to A. In contrast, when a net reinforcement generates shift in the direction of greater imbalance (i.e., lessens the information problem), it may discourage the side from risking war as it is much clearer for A what the outcome of

[3] In fact, changing certain assumptions in the model would result in the prevention of external support at the time of power mediation. See section 6.3 for the discussion of this change.

a potential battle would be.

Hypothesis 2: Net reinforcements toward greater imbalance prolong peace and net reinforcements toward greater parity destabilize peace.

Second, while power mediation "delays the inevitable" by creating unstable ceasefires, the survival of imposed ceasefires may be shaped by external support. Power mediation operates through inflating the costs of fighting. When provided during the conflict of interests stage, the sides observe great costs of fighting and derive that fighting is prohibitively costly and return to the peaceful stage for as long as pressure lasts.[4] Once the costs deflate to their previous amount, the sides will again find themselves dissatisfied with the status quo if no redistribution of the disputed good takes place. The reason why redistribution may take place in my model is the delayed materialization of support provided during the conflict of interests stage.

The effect of power mediation lasts as long as the mediator is willing to generate the barriers to fighting, whether it is threats of sanctions, potential of intervention, physical barriers to fighting, or financial incentives not to fight. In the extended model, the non-waning pressure that generates longer peace spells is the result of repeated instances of provision of power mediation.[5][6]

[4] This means that one's assumption about the duration of mediation's effect is critical to the conclusions one could derive about the effectiveness of power mediation; see section 6.3 for the discussion of possible approaches to this problem.

[5] In reality, the same mediator could also exert pressure repeatedly and extend peace. Although my model does not explicitly prohibit the same mediator from exerting repeated pressure, it is the order in which third parties decide whether to mediate that makes the same agent unlikely to be randomly chosen to mediate repeatedly.

[6] My model is consistent with current research on third parties prolonging peace through creating additional barriers to fighting or incentives to avoid fighting. For an example of barriers to fighting see work by Fortna (2004,

The probability of imposed ceasefire failure may be lowered if external support leads to redistribution of the disputed good in the dyad, such that the potential benefits from fighting a stronger as opposed to a weaker type of B decrease in the eyes of A and she becomes less likely to risk war (the result of a marginally greater preponderance). Since the redistribution of the disputed good occurs, A gives up a greater proportion of the prize than she would have otherwise. Thus, my model highlights that peaceful redistribution of the prize is more likely to take place than a military contest, when external support alleviates the information problem.

Hypothesis 3: Net reinforcements toward greater imbalance after an imposed ceasefire decrease the likelihood of armed conflict and increase the likelihood of peaceful renegotiation of the status quo.

In contrast, external support may also undermine an unstable ceasefire imposed by power mediators even further. Power mediation may be followed by a longer war than the war that could occur without support, if third parties exacerbate the information problem and make the potential benefits of fighting a weaker B as opposed to a stronger B more attractive (an exacerbated information problem is the result of marginally greater parity).

2008) on peacekeeping. For an example of third parties providing financial incentives to avoid fighting, see Arena and Pechenkina (N.d.) on US aid to Israel and Egypt.

Hypothesis 4: Net reinforcements in the direction of greater parity after an imposed ceasefire increases the likelihood of ceasefire failure.

In the sections to follow I test the former two hypotheses of the extended model, leaving the latter two hypotheses as my future contributions to the literature.

5.2 Analysis of War Duration

I now turn to evaluating *hypothesis 1* against observational data. Ideally, my data would include a sample of both inter- and intra-state conflicts with a record of the totality of military reinforcement received by each side of each conflictual dyad. Due to disparities in how scholars of international and civil conflict record external support, I follow the convention and perform my analysis employing a sample of only one "type" of war. In this section, I focus on civil conflicts, although I also expect the implications of my model to be relevant for interstate crises.

5.2.1 Intrastate crises data

The data for the analysis come from Regan (2002), which is the most comprehensive data set on third party intervention in civil conflict. Regan has recorded data on intervention events operationalized as sending troops or navy or airforce or equipment or military advisors, which is a great advantage over the Correlates of War data on joining in civil war (Sarkees and Wayman 2010).[7] Another advantage of these data is the criterion for case selection: Regan includes the total of 140 civil conflicts in 1944-1999 with at least 200 fatalities, as opposed

[7] For an example of employing the COW data, see Balch-Lindsay, Enterline and Joyce (2008).

to the standard 1,000 fatalities threshold. This is important for the study of intervention, as potential interveners may not predict perfectly how intense the fighting may become when then they decide to intervene, allowing researchers to examine a fuller portfolio of intervention options (see Gent (2008); Regan (2002) for a more detailed discussion of this point).

Dependent variable: Time until the end of violence (as coded by Regan (2002)) is the dependent variable.

Independent variables: It is not relevant to my model to focus on whether the government or the rebels receive external support. In stead, by comparing the relative size of the government and opposition's military force, one may recode the data to identify which side of the dyad is stronger/weaker at any point in time.[8] While governments in fact dominate most of the dyads in the sample, the rebels dominate 15.7% of the sample (22 out of 140 wars). The relative power is measured as the size of an actor's military force. Furthermore, 18.6% of all instances of provisions of external support to rebels occur within these cases. The average duration of a conflict is 44.14 months, while the conflicts with rebels having more military personnel than governments last on average 26.95 months, with a median of 8 months. I code REINFORCEMENT TO STRONGER SIDE as present if external support is given to the government side of the conflict and the government has a greater military force than

[8] Gent (2008) introduces the idea of measuring relative rebel capabilities by comparing how numerous the sides' military forces are. Cunningham, Gleditsch and Salehyan (2009) have introduced more detailed data on rebel capabilities, and as a future extension I will merge their capability data into Regan's data set. A different measure of relative capabilities has been introduced in DeRouen and Sobek (2004): a ratio of the size of the government's army to the total size of the population. Gent (2008) and Cunningham, Gleditsch and Salehyan (2009) point out that the total size of the population may not be a good predictor for potential opposition recruitment given the variety of factors that may impede it. All of these studies imply that government and opposition are inherently different types of actors, while in this exercise I attempt to demonstrate that bargaining dynamics between governments and oppositions conform to a general model of dyadic interaction and third-party influence.

the rebels, while the variable is coded as present if outside actors reinforce rebels and rebels dominate the government. Analogously, REINFORCEMENT TO WEAKER SIDE is coded as present if third parties provide support to a weaker government and if reinforcement is given to the rebels that control a smaller military force than the government. In the time-varying data setup, I use a decay function to reflect a waning effect of intervention over time. More specifically, intervention is coded as a dummy variable indicating whether an intervention has occurred in a given month multiplied by a decay function that decreases as time passes from the time of the most recent intervention. The presented results show effects of intervention that decay at a 8.33% per month rate, thus disappearing in 12 months. Recall that my model suggests that a net reinforcement is the measure of interest as reinforcements that take time simultaneously undermine the effect of the opposing side's support. To understand net reinforcement on each side, I interact the two major independent variables.

I use the measures of REINFORCEMENT TO STRONGER SIDE and REINFORCEMENT TO WEAKER SIDE as proxies for net external investments in the direction of greater imbalance and greater parity respectively. I imply that when a stronger side grows even stronger, the preponderance in the dyad increases, while when a weaker side is being reinforced, the gap between the sides' capabilities shrinks. These measures have obvious limitations as they do not account for the amount of received reinforcement, instead the measurements indicate instances of provisions of external support. It is conceivable that a weaker side could receive reinforcement that is disproportionately large, such that the balance would surpass parity and approach imbalance with the formerly weaker side now dominating the opponent. While such occurrences are rare, they highlight the need for more fine grained data that also indicate the amount of support given as opposed to the incidence of support.[9]

Controls: PARITY is constructed as $Parity_{d,t} = 1 - \left(\frac{MilForce_{Hi,t}}{MilForce_{Hi,t} + MilForce_{Lo,t}} \right)$, where $MilForce_{Hi,t}$ denotes the military force in month t and $MilForce_{Lo,t}$ denotes the smaller military force in that dyad. The measure varies from 0 to 0.5, with 0 indicating complete preponderance and 0.5 being close to parity. I use this variable as an identity-neutral measure of distribution of capability in the dyad, instead of the variable RELATIVE REBEL CAPABILITIES introduced in Gent (2008) that reflects the relative size of the rebel military force to the government's military force. See Figure 5.1 for the comparison of two measures.

For comparability with previous research I include the controls incorporated in the models of intervention and civil war duration by Regan (2002); Regan and Aydin (2006); Gent (2008). INTENSITY records the average monthly death rate, i.e. it is a ratio of the total battle deaths to the number of months that the conflict lasts. Higher incidence of fatalities have been previously demonstrated to shorten civil war duration.

ETHNIC/RELIGIOUS indicates whether the conflict mainly concerned ethnic or religious differences and is expected to have a negative effect on the probability of rebel victory, as previous studies have found that opposition forces are less likely to win in identity-based civil wars (DeRouen and Sobek 2004; Regan 2002; Regan and Aydin 2006; Gent 2008).

Regan (2000) has also demonstrated that the post-Cold War period has experienced more intervention events, therefore I control for that change in the international environment with a COLD WAR dummy variable.

[9] Regan (2002) codes the amount of each type of military support, however the amount of missing data in those measures is prohibitively great such that over half of cases would drop.

Finally, MOUNTAINS is a measure of mountainous terrain, as introduced in Fearon and Laitin (2003).

Method: I use semiparametric Cox event history models, Weibull parametric, and Weibull duration with selection models. I first use the Cox event history models to replicate Gent's (2008) results and demonstrate the differences in my model compared to Gent's. I also test for the violation of the proportionality of hazards assumption. Since none of the variables' effects vary over time, I do not include any interactions with the log of time (Box-Steffensmeier and Zorn 2001; Box-Steffensmeier, Reiter and Zorn 2003). As I need to account for the selection of third parties into the choice of providing support, I also use the Weibull duration with selection model (Boehmke, Morey and Shannon 2006). To demonstrate that the full effect of the Weibull model with selection, I also include the Weibull parametric and probit models to facilitate comparison between time-varying and non-time varying versions of the data.

Results: I turn now to the results of my analyses. Table 5.1 reports the results of my analysis of the probability that a civil conflict terminates. I estimate a Cox semi-parametric regression model, which is commonly used to model duration for time series cross sectional data, while also making fewer assumptions about the shape of the hazard of failure. The two models in Table 5.1 underscore the differences between exploring the effects of intervention when measured by the identity of the recipient and by the power structure in the dyad: model 1 of Table 5.1 shows the interaction and the constitutive terms of government- and rebel-biased provisions of external support and model 2 of the same Table records external support based on whether the stronger or the weaker side of the dyad receives it.

Recall that the previous research has been investigating the impact of identity-based interventions on duration until a certain outcome, therefore the comparison with Balch-Lindsay, Enterline and Joyce (2008); Gent (2008) is not straightforward. Both works have established that rebel-biased intervention shortens the time until opposition victory and civil war terminations in a negotiated settlement. Gent (2008) finds that government-biased intervention has no effect on war duration until government victory and Licht (2011) reestimates the findings in Balch-Lindsay, Enterline and Joyce (2008) to find that government-biased intervention prolongs the duration of war until government victory. Given that most wars that experience intervention receive multiple instances of intervention, it is important to separate the net effects of intervention on each side by introducing a multiplicative interaction of the two intervention types. To understand the effect of each type I use linearly combined coefficients. When estimating duration until the end of war as opposed to one side's victory over another, both types of intervention have a decreasing effect on the hazard of failure. The linearly combined coefficient for government-biased intervention is -1.2, which is statistically significant at the 0.5 statistical significance level. The linearly combined coefficient for the rebels-biased intervention is -0.73, and it is not statistically significant. Similarly, when the intervention types are recoded as support of the stronger or weaker side, both linearly combined effects indicate that both types of intervention prolong war. The effects of control variables are consistent with the published studies: wars with higher death tolls are more likely to end sooner, while ethnic/religious wars take longer to end. Also, the Cold War indicator suggests that during the Cold War civil conflicts were more likely to end sooner. A more difficult terrain is associated with longer wars (although not significant in model 1 of Table 5.1), and political regime has no effect on war duration.

While a survival model with time varying covariates allows me to explore how changes within panels over time affect war duration, it does not account for why some third parties select themselves to intervene while others stay out. To make sure that these results are not affected by the motivations of third parties to intervene in the first place, I also employ a Weibull with selection model, which is only available for a non-varying duration data setup. Therefore, I collapse the data by conflict, converting a time series cross section data set into a cross-sectional format with 140 observations, such that each observation represents a civil conflict. The independent variables are transformed into the counts of instances of reinforcement events. In other words, if actor i receives 5 instances of reinforcement, the variable will be set equal to 5. The mode of number of reinforcements is 0, the mean is 2, and the maximum is 30.

Consider Table 5.2, which demonstrates a naive Weibull parametric survival model applied to the cross-sectional version of these data (model 1 of Table 5.2), a probit model estimating what factors determine which conflicts receive intervention (model 2 of Table 5.2), and a Weibull duration model with selection that estimates duration until civil war's termination accounting for whether a civil war experienced intervention by outside actors (model 3 of Table 5.2). First, compare the coefficient estimates in model 1 of Table 5.2 to model 2 of Table 5.1: the principal difference between the two models is the format of the data (cross-sectional in Table 5.2 and time-series cross sectional in Table 5.1). While the effects of some of the control variables change when I eliminate variation within each panel (i.e., the covariates that measure Cold War and mountainous terrain are no longer statistically significant in the collapsed data), the major relationships between the net reinforcements to either side of the dyad and war duration are preserved. The linearly combined coefficients are -0.27 and -0.35 for the stronger side-biased and the weaker side-biased net reinforcements respectively. Again, when not controlling for the selection effects, both types of intervention are associated with longer civil conflicts.

In contrast, model 3 of Table 5.2 yields a linearly combined coefficient of 0.066 for the stronger side-biased net reinforcement. Given that the variable REINFORCEMENT TO STRONGER SIDE is measured as the sum of instances of provisions of support, the substantive effect of one instance of support would be $100 \times [exp(0.066) - 1] = 6.823\%$ increase in the probability of war termination compared to those case that did not experience a net reinforcement on the side of the stronger actor. This result is consistent with my model which yields that reinforcements in the direction of greater imbalance should help wars end sooner due to the lessening of the information problem. The linearly combined coefficient for the weaker side-biased reinforcement remains negative, -0.016, even after controlling for the selection causes of provisions of external support.

Consider a selection with duration model (shown as model 4 of Table 5.2) that includes identity-based indicators of external reinforcement. While the directions of the effects are analogous to model 3 of Table 5.2, the relative sizes of reinforcement coefficients differ significantly depending on whether one employs identity-based or power structure-based measures of net reinforcement. The linearly combined coefficients in model 4 of Table 5.2 yield that support to government shortens war, albeit the effect is not statistically significant, while external reinforcements of rebels also lead to shorter wars, and the effect is statistically significant.

Overall, the results suggest a few takeaways. First, my taking into account of the selection stage that leads to third-party reinforcement generates substantially different results with respect to whether reinforcement prolongs war or shortens it. Specifically, the effects of identity-based and power-based measures of intervention appear to be analogous when selection is not taken into account, yet these measures generate different results when selection is accounted for. Second, Gent (2008) acknowledges that sometimes rebels dominate the government but states the percentage of such cases is insignificant. I demonstrate that recoding support indicators based on power structure in the dyad generates results quite different from identity-based indicators in the full information duration model. These results are consistent with the general bargaining model of conflict management, which is suggestive of civil wars exhibiting analogous bargaining dynamics with other types of conflicts.

5.3 Analysis of Peace Duration

5.3.1 Peace duration after interstate wars

I now turn to the evaluation of *hypothesis 2* against a sample of ceasefires signed by recognized governments.

Data and method: The data are time series of multiple sections and come from Lo, Hashimoto and Reiter (2008); Ghosn, Palmer and Bremer (2004), and the USAID website. Each section starts with the end of an interstate war. The end of each section is determined by the dependent variable: (i) when using NEW WAR as an indicator of ceasefire failure, each section ends when the new war starts,[10] (ii) when using VIOLENT MID, each section ends when the new violent militarized interstate dispute starts (hostility level = 4-5).

[10] Lo, Hashimoto and Reiter (2008) include the interstate wars that reached 1,000 battle deaths as recognized by the Correlates of War (COW) Project (Sarkees 2000).

The data were introduced by Lo, Hashimoto and Reiter (2008), who document 186 ceasefires in 1914-2001 (a significant expansion from Fortna's data set of 48 ceasefires in 1946-1998). Out of 186 ceasefires, 54 experience a new war; in other words, war resumes in about 29% of ceasefires. The data include more than 2,673 ceasefire years. The shortest ceasefire lasts 5 days and the longest failed ceasefire lasts almost 58 years. The longest ceasefire that never failed as of 2001 is 87 years.[11] The unit of analysis is the dyad-ceasefire year. Multilateral wars are disaggregated into dyads.[12] Each subject is a ceasefire-dyad, and the dependent variable is the length of time until that dyad fights a new war.

Independent Variables: In the previous section the reinforcement data available to me were measured in various units, therefore it would have been meaningless to compare the amounts provided; I focused on the counts of instances of provisions of support. Here, all reinforcements are measured in 2005 constant US dollars, which means that net reinforcement is a meaningful concept for these data. Ideally, net reinforcement would be measured as the difference between the total military aid received by A and B from every other actor in the international system. For this study, I use a proxy measure of net reinforcement, represented as the difference between A's military aid and B's military aid received from the US.[13] My model assumes that third-party reinforcements are an instrument for achieving a certain policy outcome. The US does not distribute aid randomly across the world's governments, rather aid is an expression of the American government's attempts to achieve

[11] Whether a new war is a failure of ceasefire established after the previous war that took place 58 years ago is an important theoretical question. Here, I follow the literature and do not use any temporal cut points for ceasefire duration. This theoretical question is one of the directions that future research should address.

[12] See (Lo, Hashimoto and Reiter 2008, p.722-723) for a detailed treatment of disaggregation of multilateral wars.

[13] I also use a value of US military aid lagged by one year; the results are not affected by that change.

particular distributional outcomes within the world's pairs of actors. While the US is the largest provider of military aid in the world, which implies that this measure accounts for a significant amount of military aid, it does not represent the full amount and therefore needs to be interpreted narrowly. This is an example of a single third party's (albeit the most powerful one's) attempts to change policy, not as the total amount of attempts across the world to shape distributional outcomes among actors. The variable is measured in total 2005 US dollars. I also include the difference in reinforcement squared, as the overall fit of the model is improved by this inclusion. Importantly, the squared term does not alter the direction of the effect, therefore this inclusion does not contradict the model's implications.

The relative distribution of capabilities in the dyad is measured as the proportion of dyad's resources controlled by side A, such that 0.0 means B controls 100% of the dyad's resources, 0.5 implies that the sides are at exact parity, and 1.0 means A possesses all resources in the dyad. The average dyad in these data has A controlling at least 60% of the capabilities, with the maximum value of 99% and the minimum amount of resources possessed by side A is 1.08%. Resources are measured using the capabilities measure from COW Composite Index of National Capability (CINC) data set (Singer, Bremer and Stuckey 1972). Figure 5.2 charts the density distribution of this ratio. Importantly, this measure changes yearly, therefore, any shifts in capabilities that may occur for the reasons other than third party influence will be incorporated in this measure (e.g. natural disasters, regional economic crises, discovery of new weapons systems).

For this analysis, I use the interaction term of net difference in reinforcement and ratio of capabilities as my major independent variable as explained in Brambor, Clark and Golder (2006).

Controls: TIE was originally introduced in Fortna (2003). This is a binary variable that indicates whether a war ended in a military stalemate (1) or a victory for one side (0). It captures whether at the end of the war the belligerents have contradictory expectations about the outcome of a next potential battle, which is the case if a war ends in stalemate. A measure of military stalemate has been a standard control variable in peace duration studies.

INDEX OF AGREEMENT STRENGTH captures how many various provisions or characteristics designed to stabilize a ceasefire an agreement has (Fortna 2003, 2004). The theoretical expectation is that the more such provisions or characteristics are included, the stronger the agreement. This measure is composed of equally weighted provisions: the withdrawal of forces, demilitarized zones, arms control, peacekeeping, and internal control. In addition, equally weighted characteristics of an agreement are included: external involvement (whether a third party state mediated a ceasefire or provided guarantees of peace), confidence building measures (whether the sides exchanged information on militarily critical aspects of their defense systems, e.g., mines), and dispute resolution (whether a third party provides dispute resolution channels or the belligerents established a joint commission to resolve issues). Finally, INDEX OF AGREEMENT STRENGTH includes a measure of detail, which is built on paragraph count and approximates how specific an agreement is (Fortna 2002).

Following the literature on peace stability, I also include controls for the characteristics of war employed in previous studies. COST OF WAR is a continuous variable that records the log of battle deaths generated in a war. HISTORY OF CONFLICT is a ratio of the number of serious disputes the dyad experienced to the duration of the dyad's existence. EXISTENCE AT STAKE is a binary measure that reflects whether one side's very existence was threatened by the war. CONTIGUITY is a dummy variable that indicates whether the dyad is contiguous by land or separated by less than 150 miles of water.

Method: I use semiparametric Cox event history and probit regressions. The former model helps determine whether any of the variables fail tests for proportional hazards and need to be corrected by introducing the interaction terms of those covariates with the log of time (Box-Steffensmeier and Zorn 2001; Box-Steffensmeier, Reiter and Zorn 2003). In the models presented, none of the variables violates this assumption.[14] The probit regression models are used with a time variable, whose square and cubic interactions account for time dependency just like survival models do (Beck, Katz and Tucker (1998); Carter and Signorino (2010)).

Results: Consider model 1 of Table 5.3 that reports three models determining the recurrence of new war recurrence, of new violent MID (model 2 of Table 5.3), and new nonviolent MID recurrence (model 3 of Table 5.3). I estimate a probit regression, which is commonly used to model binary outcome dependent variables for time series cross sectional data. Overall, the results are consistent for modeling the effects of the specified independent interacted covariates on the probability of *New War* and *New Violent MID* and have no effect on *New Nonviolent* $MID_{d,t}$. In other words, the effects are consistent across two different measures of material conflict, yet inconsistent between the measures of material and verbal conflict. While the net reinforcement measure does not change direction, interpreting squared interaction terms is not straightforward without visualization. To facilitate interpretation of

[14] The Cox models are omitted; the effects are very similar to the ones shown in the probit regression models.

the measures of interest, consider Figure 5.3, which charts the predicted probabilities of war recurrence (see y axes) as a function of net reinforcement and balance of capabilities in the dyad. The x axes show the difference in military aid received by A compared to B measured in millions of 2005 constant US dollars, such that 0 denotes no difference in the amount of US military aid received by A and B, negative numbers indicate B receiving relatively more aid and positive numbers show A receiving relatively more military reinforcement from the US.

Both panel (a) and panel (b) of Figure 5.3 show the predicted probability of war recurrence for a dyad in which A dominates the opponent by controlling 90% of dyad's resources in blue (solid lines chart predicted probability and dash lines denote 95% confidence intervals). When A dominates her opponent, additional reinforcement will make A even more powerful. The extended model suggests that greater amounts of reinforcement in the direction of greater imbalance eliminate the information problem and the sides are less likely to risk military contest. In addition, my model suggests that peaceful transfers of resources would occur under such circumstances. Unfortunately, due to data limitations, I am not able to measure peaceful renegotiations of the status quo. The graphs demonstrate that even when B is being reinforced by 2.5 billion dollars more compared to A: the probability of military contest rises, however, the confidence intervals are wide and include zero, therefore the variance increases when A is being reinforced, but the effect is not statistically significant.

The difference between the two panels lies in what scenario I compare the dyad in which A controls 90% of resources to. Panel (a) of Figure 5.3 shows in red the predicted probability of war recurrence for a hypothetical dyad in which A controls 10% of resources. As the model suggests, reinforcements that generate shifts in the direction of greater parity (which is the case here) increase the probability of ceasefire failure. This is consistent with the extended model.

Panel (b) of Figure 5.3 charts in red the predicted probability of war recurrence for a hypothetical pair of governments, in which A controls half of the dyad's total capability. Given that growing from half implies that the sides move toward greater imbalance, one would expect a pacifying effect of such net reinforcement, exactly as shown in panel (b).

Consider Figure 5.4 that presents the same comparison of cases for another dependent variable: occurrence of violent MID. Panel (a) demonstrates that the effect of net reinforcement is weaker when it comes to lower level violence: it takes 5 billion in reinforcement to generate a statistically significant increasing effect on ceasefire failure (when failure is measured as violent MID), as opposed to 500 million when failure is recorded as new war. This is an interesting result that requires further investigation. Furthermore, when one compares the effect of net reinforcement for a weak A on the probability of nonviolent MID, panel (a) in Figure 5.5 shows that there exists no statistically significant relationship between net reinforcement and verbal conflict among recognized states of the international system.

In short, the major substantive takeaway from this analysis is that the observational data on interstate ceasefires reveal that the extended model's implications are consistent with the behavioral patterns of war recurrence. The patterns of violent behavior that does not reach the level of war is consistent with the extended model, but the substantive effect of net reinforcement on violent MID occurrence is not as great as it is on war recurrence. Finally, the model's implications are not consistent with occurrence of verbal conflict in the international system (as recorded in this sample). Although my model does yield the implication that peaceful transfers of resources are more likely to take place under the net reinforcement in the direction of greater imbalance, it is not obvious whether nonviolent MIDs may be viewed as indicators of peaceful renegotiations of status quo. Rather, a nonviolent MID may or may not involve some degree of renegotiation, but nothing in the coding rules of the data set suggests that renegotiation was taken into account for this variable (Ghosn, Palmer and Bremer 2004).

The effects of all the control variables are consistent with prior research on the subject with the exception of CONTIGUITY generating a statistically negligent impact on the probability of conflict recurrence. I have examined whether the measure violates the assumption of proportionality of hazards. Although the tests the proportional-hazards assumption on the basis of Schoenfeld residuals does not yield that CONTIGUITY violates the assumption, when corrected by interacting contiguity with the log of time the initial destabilizing effect becomes positive and statistically significant. It does not affect the rest of the covariates, therefore I omit the interaction.

5.4 Conclusion

In this chapter I have attempted to test two out of four major hypotheses developed in the extended agent-based model of chapter 4. First, I tested whether relative reinforcements that shifts the balance of capabilities in the direction of imbalance within a civil war dyad precipitates war termination (hypothesis 1). To approximate the direction of relative reinforcement I code whether reinforcement is provided to a stronger as opposed to a weaker side. Such a coding scheme violates the convention in civil war research, whose focus has been on the effects of reinforcements that are coded based on the identities of the recipients. I demonstrate that although government actors in civil war dyads are overwhelmingly more likely to be stronger than the opposition, there is enough variation in the how capabilities are distributed within civil war dyads to make my coding scheme meaningful. Once I account for the selection effects related to third parties choosing to reinforce some civil war actors and not others, I demonstrate that the effect of reinforcement of the stronger side is associated with shorter wars, while reinforcements of the weaker participant in the civil war is associated with a longer civil war. Tellingly, when I replicate the same duration with a selection model, the conventional coding scheme yields that both reinforcements of government and rebels have a shortening effect on war duration. I believe my analyses are suggestive that a more general coding scheme based on the power structure of the civil war dyad provides a greater insight into the bargaining process of civil wars: it is analogous to other types of conflict. An important extension to this analysis would be an examination of power-based intervention in international crises. The convention in the international conflict literature has been to examine intervention based on whether it reinforces the initiator or the target in the conflict, which is a problematic distinction and is not general enough to apply to other types of conflict.

Second, I have tested whether relative reinforcements after a ceasefire is signed have a pacifying or destabilizing effect on ceasefire survival (hypothesis 2). My model suggests that conflicts of interests recur regularly in the international system and third-party reinforcement may encourage or discourage the sides from escalating a conflict of interests to a military contest. I use a sample of ceasefires established after interstate wars of the 20th century and test whether power structure of the dyad combined with relative reinforcements is related to conflict behavior. I find that my model's implication is consistent with conflicts of higher levels violence however the logic of relative reinforcements does not apply to nonviolent patterns of conflict in the international system. Given that in the model a rejection of A's offer leads to a battle, this lack of relationship between nonviolent conflict and patterns of reinforcements does not undermine the model.

Fig. 5.1. Histograms of key measures of rebel capability for intrastate crises sample

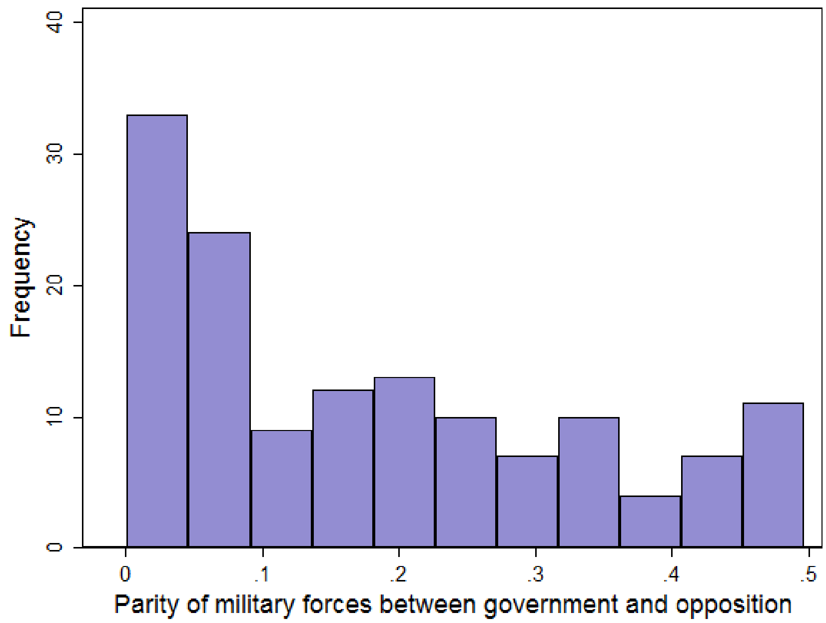

(a) $Parity_{d,t} = 1 - \left(\frac{MilForce_{Hi,t}}{MilForce_{Hi,t} + MilForce_{Lo,t}} \right)$, where $MilForce_{Hi,t}$ denotes the military force in month t and $MilForce_{Lo,t}$ denotes the smaller military force in that dyad, such that 0 = complete domination by one side, and 0.5 = complete balance of armed forces

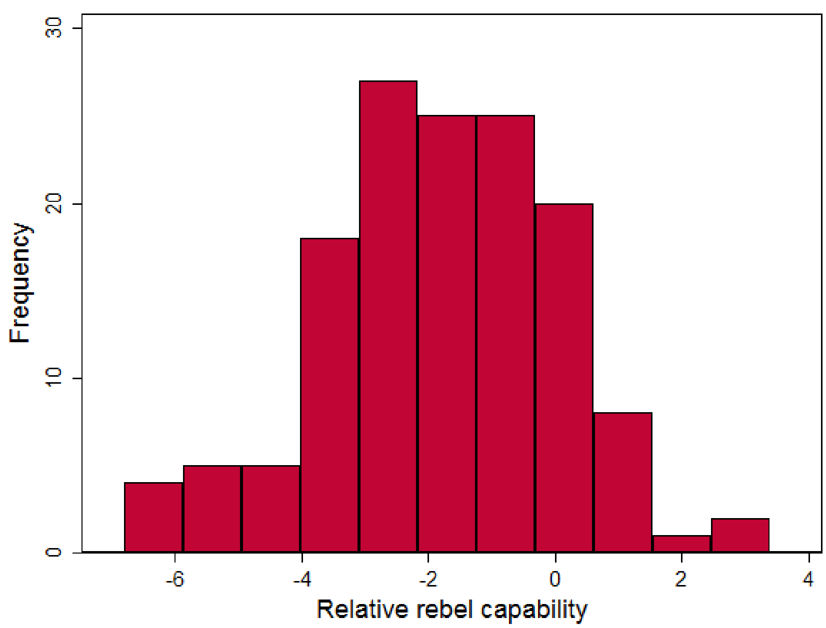

(b) Relative size of rebel military force: $ln(\frac{RebelTroops}{RebelTroops+GovtArmy})$

Fig. 5.2. Histograms of key independent variables for the interstate ceasefires sample

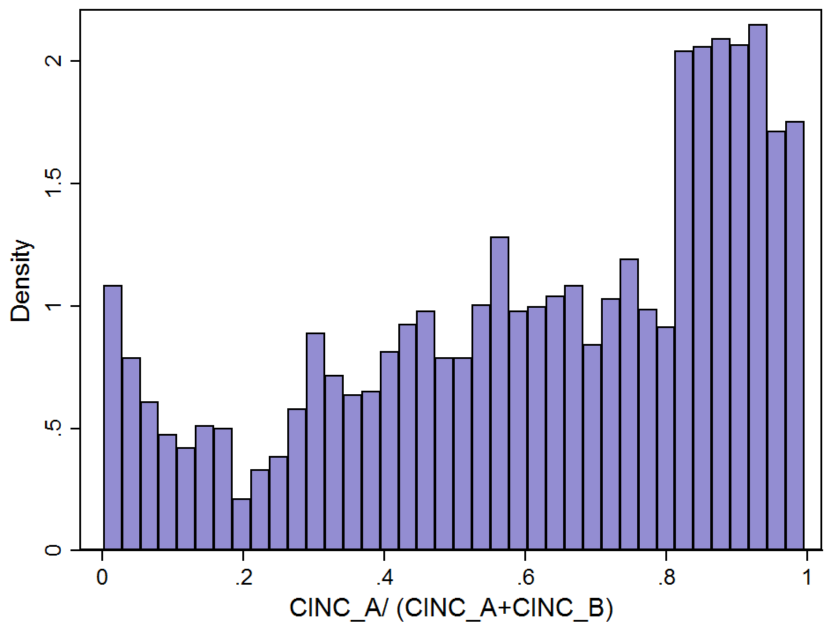

(a) A proportion of the dyad's resources controlled by A

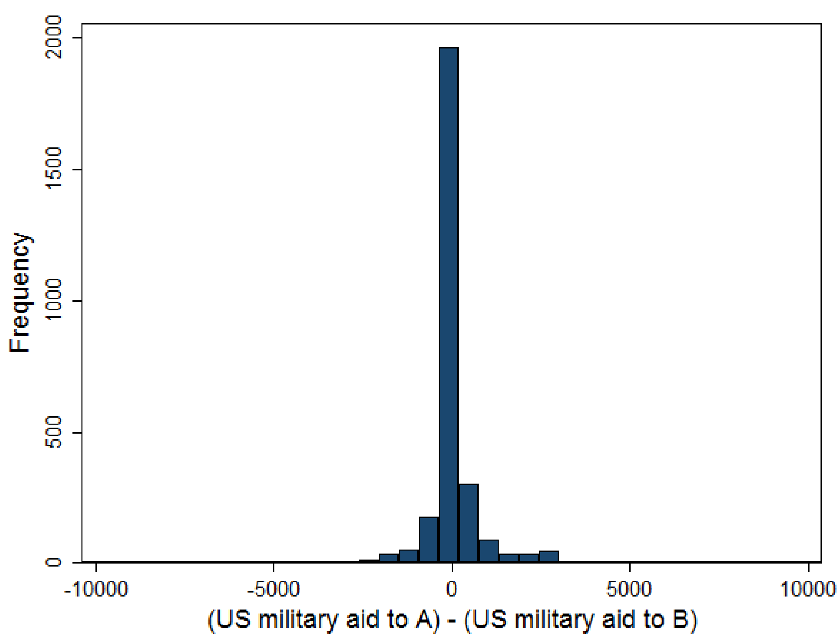

(b) A difference between the amount of US military aid received by A and B

Fig. 5.3. The effect of external support on the probability of war recurrence

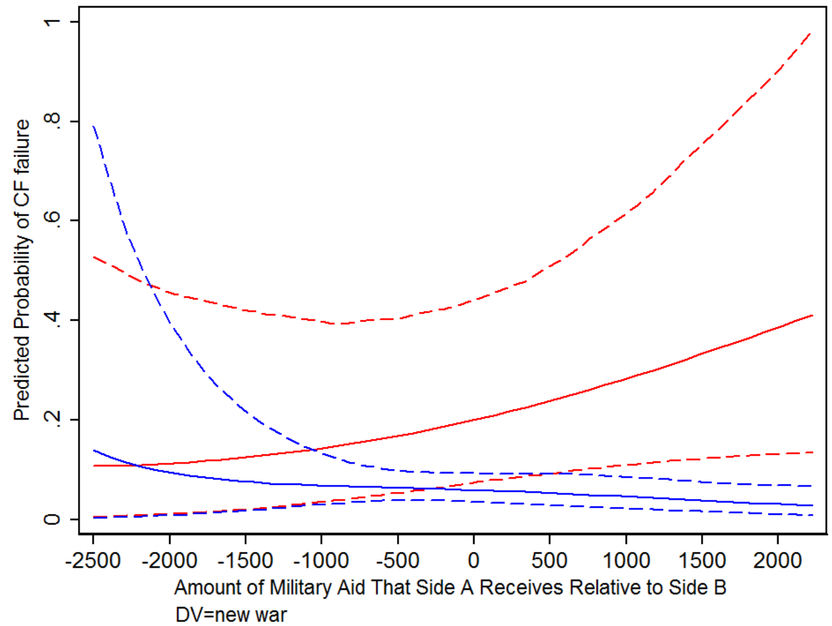

(a) Red line: A controls 10% of the dyad's capability, Blue line: A controls 90% of the dyad's resources

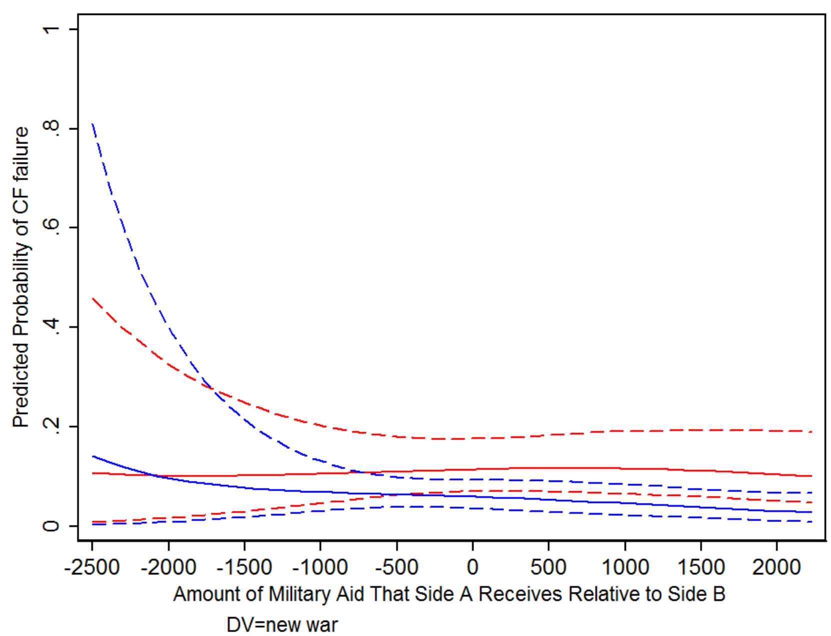

(b) Red line: A controls 50% of the dyad's capability, Blue line: A controls 90% of the dyad's

Note: The dependent variable is new war. Military aid difference is measured in millions of 2005 constant US dollars. Positive values on the x axis represent side A receiving more military aid than side B, while the value of 0 implies that both sides have received equal amounts of military aid from the US and negative values represent side B receiving more military aid than A. The graph shows that for those cases in which A controls 90% of the dyad's resources (blue lines), reinforcements of A are pacifying, while reinforcements of B encourage war. For those cases in which A controls 10% of the dyad's resources, reinforcements of A increase the probability of ceasefire failure (red lines, panel A), while when A controls half of the dyad's resources, the effect is not statistically

Fig. 5.4. The effect of external support on the probability of ceasefire ending with a violent MID, levels 4-5

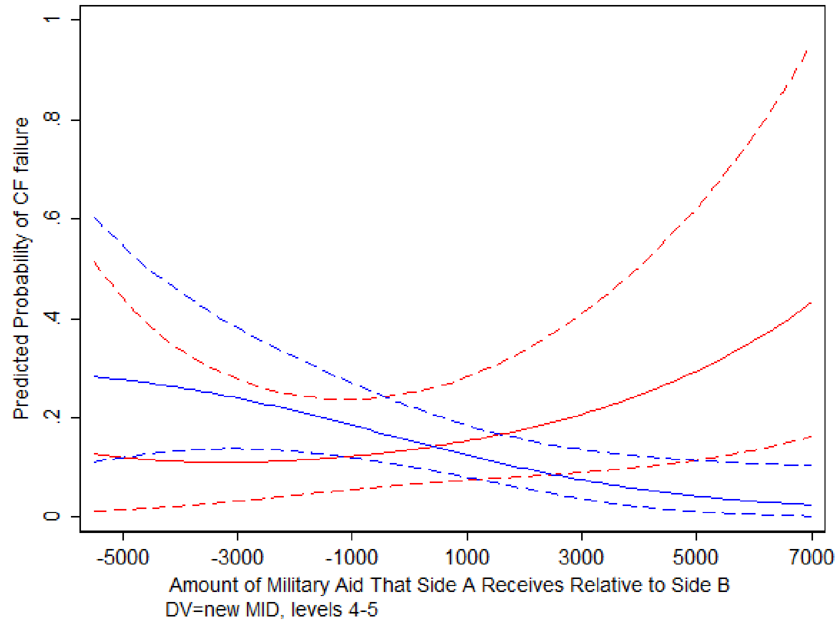

(a) Red line: Acontrols 10% of the dyad's capability, Blue line: A controls 90% of the dyad's resources

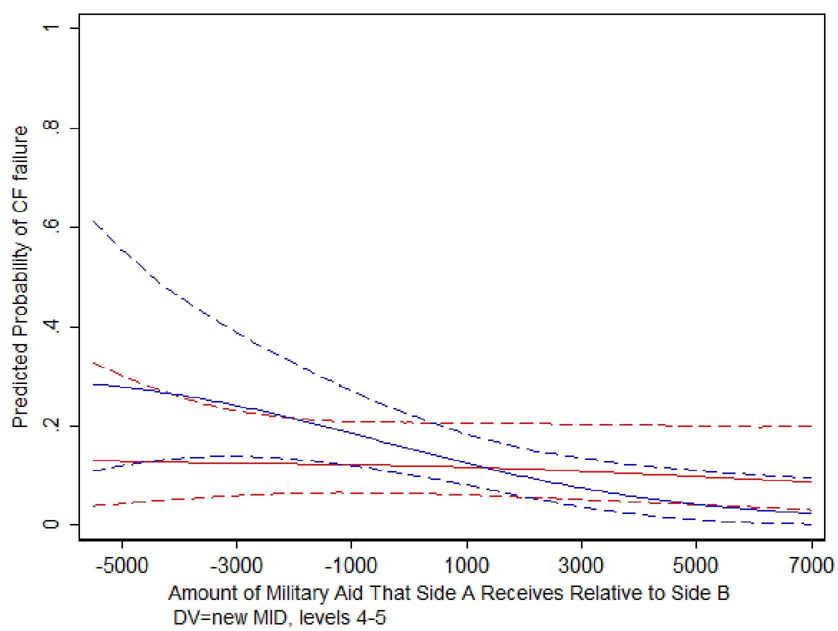

(b) Red line: A controls 50% of the dyad's capability, Blue line: A controls 90% of the dyad's resources

Note: The dependent variable is MID, levels 3-5. Military aid difference is measured in millions of 2005 constant US dollars. Positive values on the x axis represent side A receiving more military aid than side B, while the value of 0 implies that both sides have received equal amounts of military aid from the US and negative values represent side B receiving more military aid than A.

Fig. 5.5. The effect of external support on the probability of ceasefire ending with a nonviolent MID

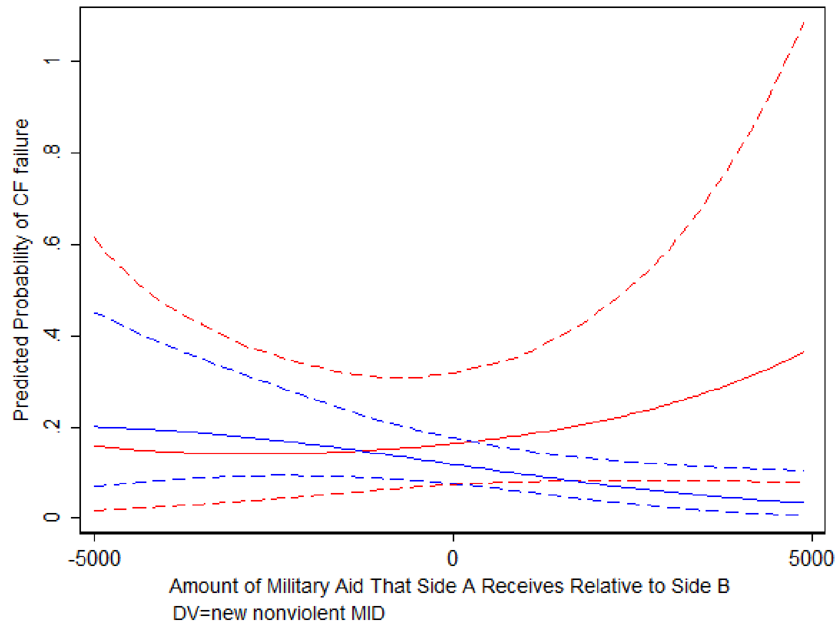

(a) Red line: A controls 10% of the dyad's capability, Blue line: A controls 90% of the dyad's resources

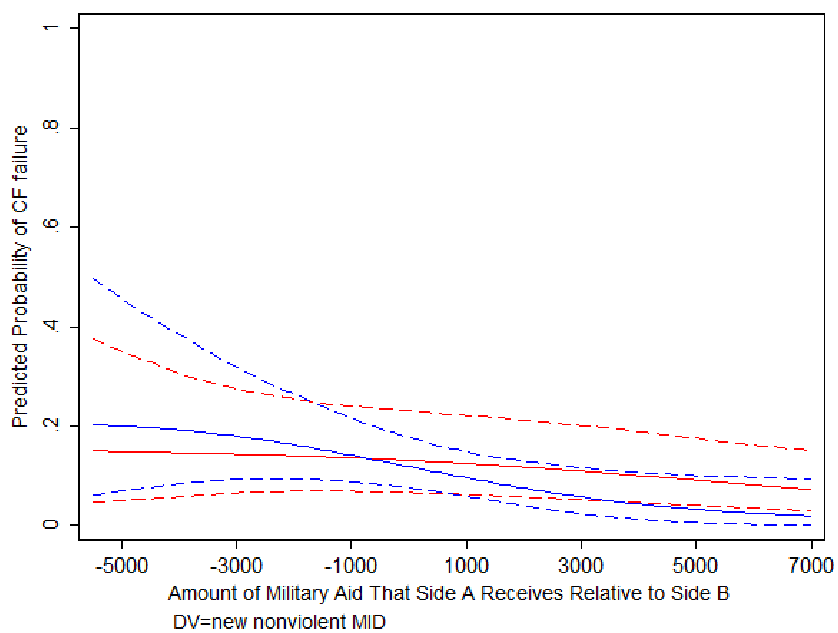

(b) Red line: A controls 50% of the dyad's capability, Blue line: A controls 90% of the dyad's resources

Note: The dependent variable is MID, levels 3-5. Military aid difference is measured in millions of 2005 constant US dollars. Positive values on the x axis represent side A receiving more military aid than side B, while the value of 0 implies that both sides have received equal amounts of military aid from the US and negative values represent side B receiving more military aid than A.

Table 5.1. Cox semi-parametric regressions of civil conflict duration

Covariates	1 Identities-based Approach	2 Power Structure Approach
. Government- $\times$ Rebel-biased Intervention	-1.59* (.73)	
Government-biased	.30 (.32)	
Rebel-biased	.86* (.38)	
Stronger side $\times$ Weaker side-biased Intervention		-1.39* (.70)
Stronger side		.50‡ (.32)
Weaker side		.69* (.39)
Intensity	.04**(.01)	.04** (.01)
Ethnic/Religious	-.31*(.19)	-.34* (.20)
Cold War	1.45** (.55)	1.42** (.50)
Mountains	-.12 (.08)	-.14* (.08)
Democracy	.002 (.02)	.001 (.01)
$Pr > \chi^2$	0.000	0.000
Log pseudolikelihood	-450.23	-442.55
N obs	11,542	11,542
N subjects	139	139

Note: $** < .05$, $* :< .1$ in a two-tail test. ‡ < .1 in a one-tail test. Robust standard errors are in parentheses and are reported clustered on conflicts. The coefficients in the duration equations are reported as influences on the hazard of failure. The negative coefficient has a decreasing effect on the hazard of failure, and thus a prolonging effect on civil war duration. The positive coefficient has an increasing effect on the hazard and thus a shortening effect on war duration. Linear combinations of interacted variables: Government-biased: -1.2*(0.66); Rebels-biased: -.73 (0.56); Stronger side-biased: -0.89‡ (0.61); Weaker side-biased: -0.69 (0.56).

Table 5.2. Controlling for the unobservable causes that lead to intervention in support of rebels in civil wars. Duration until war termination.

Covariates	1 Naive Weibull *Time until war end*	2 Probit DV=Any Int	3 Weibull w/Sel select= Any Int *Time until war end*	4 Weibull w/Sel select= Any Int *Time until war end*
Selection Eq				
Parity		.79 (.75)	-1.78 (1.52)	.80 (1.17)
Intensity		.01 (.01)	-.005 (.01)	.01 (.02)
Ethnic/Religious		-.23 (.24)	-.26 (.30)	.14* (.07)
Neighbors		.11* (.05)	.17** (.07)	-.35 (.30)
Cold War		.29 (.26)	.41 (.33)	.34 (.34)
Constant		-.31 (.36)	.02 (.44)	.02 (.43)
Duration Eq				
Government- × Rebel-biased Intervention				-.012 (.03)
Government-biased				.04 (.136)
Rebel-biased				.06* (.04)
Stronger × Weaker side-biased Intervention	-.40* (.21)		-.03* (.002)	
Stronger side-biased Intervention	.13** (.03)		.096** (.042)	
Weaker side-biased Intervention	.05‡ (.03)		.014‡ (.011)	
Intensity	.04** (.01)		-.04** (.01)	-.05** (.01)
Ethnic/Religious	-.33‡ (.21)		.04 (.31)	-.15 (.29)
Cold War	.29 (.27)		.03 (.32)	-.33 (.39)
Mountains	-.07 (.09)		.18‡ (.12)	.16 (.13)
Democracy	.006 (.02)		-.03 (.03)	-.03 (.03)
Constant	-3.08** (.44)		3.95** (.44)	3.97** (.55)
Error correlation *Rho*			-.25** (.002)	-.24 (.28)
$Pr > \chi^2$	.0001	0.005	0.0004	0.0045
Log pseudolikelihood	190.27	-87.04	-430.69	-423.38
N (uncensored)	111	140	140 (87)	140(87)

Note:$**$ $< .05$, $*$ $:< .1$ in a two-tail test. ‡ $< .1$ in a one-tail test. Robust standard errors are in parentheses and are reported clustered on conflicts. The coefficients in the duration equations are reported as influences on the hazard of failure. The negative coefficient has a decreasing effect on the hazard of failure, and thus a prolonging effect on civil war duration. The positive coefficient has an increasing effect on the hazard and thus a shortening effect on war duration. linearly combined coefficients for the interacted variables include: in Model 1, stronger side-biased reinforcement: -.27‡ (.19), weaker side-biased: -.35** (.10); in Model 3, stronger side-biased reinforcement: .066‡ (.044); weaker side-biased reinforcement: -.016‡ (.01). In contrast, model 4 yields linearly combined coefficients: government biased intervention: 0.028 (.106), rebel-biased intervention: .048* (0.01).

Table 5.3: Probit regressions of the probability of cease-fire ending with a *new war*, COW-identified observations, 1914-2001

	1		**2**		**3**	
	New war		**Violent MID**		**Nonviolent MID**	
Covariate	β	SE	β	SE	β	SE
Net Reinforcement	-0.0003*	(0.0001)	-.0002*	(.0001)	-.0002	(.009)
Net Reinforcement^2	-9.74e^{-05}	(1.72e^{-07})	-2.47e^{-08}	(1.68e^{-08})	-2.094e^{-08}	(1.474e^{-08})
Capability Ratio ×	-.02	(.43)	-.38	(.36)	.27	(.33)
Net Reinforcement × Capability Ratio	0.0006*	(0.0002)	.0003*	(.0001)	.0003	(.002)
Net Reinforcement^2 × Capability Ratio	6.03e^{-05}*	(1.86e^{-05})	3.04e^{-08}*	(2.24e^{-08}s)	2.72e^{-08}	(1.93e^{-08})
Battletide	-.15	(.37)	.15	(.36)	-.16	(.41)
Tie	.30*	(.26)	.32*	(.26)	.32*	(.26)
Agreement Strength	0.02	(.06)	.04	(.39)	.05	(.04)
Ln(Deaths)	-.12**	(.05)	-.007	(.03)	.02	(.03)
History of Conflict	.86**	(.24)	1.07**	(.14)	1.06**	(.15)
Existence at Stake	.24*	(.18)	.02	(.14)	.08	(.15)
Contiguity	.14	(.19)	.12	(.24)	.87	(.75)
Time	-.05	(.04)	-.06**	(.02)	-.02	(.04)
Time^2	0.001	0.001	.0009	(.006)	2.06e^{-02}	(.001)
Time^3	-1.19e^{-05}	(1.09e^{-06})	-4.61e^{-06}	(4.15e^{-06})	-5.11e^{-07}	(1e^{-03})
Constant	-1.64**	(.52)	-1.55**	(.35)	-1.76**	(.56)
Log pseudolikelihood	-93.18		-389.67		-497.94	
Pseudo R^2	0.21		0.26		0.24	
N obs			2,779			
N groups (war)			98			

Note: $** < .05$, $* :< .1$ in a two-tail test. For the definition of cluster identifier "war" see the codebook for Lo, Hashimoto and Reiter (2008). The negative coefficient has a decreasing effect on the probability of new conflict (prolongs peace duration). The positive coefficient has an increasing effect on the probability of new conflict (shortens peace duration).

Chapter 6

Conclusion

6.1 Summary of Contributions

My dissertation contributes to the literature on conflict management by offering a general model of third-party influence on conflict onset, duration, and recurrence. In addition, I introduce a new methodological tool to make the evaluation of simulation models more efficient. I have accomplished three major things:

- Developed and analyzed an agent-based model of conflict management in two stages: an unrealistic world of dyadic interactions with no third party influence and a system in which outside actors are able to influence bargaining dynamics of the belligerents,
- Tested two out of four major hypotheses derived from the agent-based model against observational data,
- Developed and tested an automated genetic search software application to conduct searches of parameter spaces in simulation models systematically, yet more efficiently.

In chapter 2, I rely on the bargaining literature to simulate a system of peaceful, potentially conflictual, and conflictual interactions between government agents. I assume that uncertainty is the major driving force of conflict. Side A of each dyad has to evaluate which of the three possible types of B, her opponent is. The causes of conflict onset and conflict continuation are the same. First, consistent with the previous bargaining results, the model yields that A is more likely to risk war (i.e., either risk the first battle or the second battle) if the sides' capabilities are close to parity. This is explained by the fact that A's probability of defeating a weaker as opposed to a stronger opponent grows with parity. Second, A is more likely to risk war when fighting is cheap.

The benchmark model represents a stylized version of the world in which Edward Luttwak's (1999) plea to "give war a chance" is answered. The condition of global peace, i.e., no ceasefires failing at any point in time, is achievable in the null model, due to the absence of third parties (or any other sources of changes to capabilities). Since no external actors alter the distribution of capabilities or the costs of fighting in the middle of an ongoing conflict, the belligerents settle once no uncertainty remains about their relative military power. Therefore, the null model establishes that war is an effective instrument for solving conflicts when no third parties interfere and no random shocks to capabilities occur.

The extended version of the model, introduced in chapter 4, portrays a more realistic international system: third parties shape dyadic bargaining through power mediation or reinforcement. Power mediators inflate the costs of fighting and force the sides to settle, while external supporters shift capabilities in the dyad such in the direction of their preferred policy outcome. Since power mediation creates barriers to fighting instead of addressing the underlying information problem in the warring dyad, imposed ceasefires are unstable.[1] A novel idea that emerges from my model is that those dyads that find themselves dissatisfied with the status quo (after the effect of power mediation ends) may either avoid violence due to external support or fight a longer war due to external support.

[1] This result is consistent with an emerging literature on power mediation, see Beardsley (2008, 2011); Werner and Yuen (2005).

When outside reinforcement lessens the information problem by shifting the division of power in the direction of imbalance, the benefits of fighting a stronger type of B as opposed to a weaker type of B appear less attractive to A, therefore, she is less likely to risk war and the dyad may avoid an armed conflict altogether by transferring resources peacefully in accordance with the true division of power. This idea is expressed as hypothesis 3 in chapter 5: *net reinforcements toward greater imbalance after an imposed ceasefire decrease the likelihood of armed conflict and increase the likelihood of peaceful renegotiation of the status quo.*

Given that the causes of conflict onset, recurrence, and continuation are the same in the model (parity that leads to greater uncertainty and lower costs of fighting), external support that shifts the power structure in the direction of greater imbalance (and is provided in the middle of fighting, i.e., after the first battle in the model), shortens war duration. This is summarized as part 1 of hypothesis 1 in chapter 5: *net reinforcements toward greater imbalance shorten war.*

In contrast, external support generates conflict recurrence or extends an ongoing conflict when a third party's policy preference motivates him to shift the balance of capabilities such that the information problem becomes exacerbated in the bargaining dyad. If reinforcement changes the balance of capability in the dyad toward greater parity, then the difference between A's probability of defeating two different types of B grows and A is more likely to risk war, because the benefits of defeating a stronger type of B as opposed to a weaker type of B increase. These ideas are summarized in chapter 5 as:

- hypothesis 4: *net reinforcements in the direction of greater parity after an imposed ceasefire increases an a priori high likelihood of an imposed ceasefire failure,*

- hypothesis 2: *net reinforcements toward greater imbalance prolong peace and net reinforcements toward greater parity destabilize peace*, and

- part 2 of hypothesis 1 *net reinforcements toward greater parity prolong war.*

Two reasons prevent the emergence of global peace in the extended model. First, third parties provide power mediation which leads to a "restart" of the same conflict in the future (although external support sometimes creates imbalance which may prevent an armed conflict). Second, third parties also play the roles of A or B in other dyads. Therefore, when they pay the cost of reinforcement or power mediation, the distribution of power changes in their own dyads as well. The global amount of resources decreases over time in the system due to power mediation being costly. Support ensures the transfer of resources from one agent to another. The two behaviors lead to regular conflict recurrence, because of regular changes in actors' satisfaction with the status quo.

I test two of the model's implications against observational data in chapter 5. Using a sample of civil conflicts, I test whether net reinforcements in the direction of imbalance shorten conflict, while net reinforcements toward parity prolong it (hypothesis 1). To do so, I recode external support of the stronger side in the dyad to approximate shifts toward imbalance, and support to weaker side as shifts toward parity. I demonstrate support for the agent-based model's hypothesis through my analysis of reinforcement based on the power structure of the civil war dyad as opposed to the conventional identity-based codings of support (e.g., support to the opposition or government side). This analysis suggests that civil war bargaining conforms to the general framework of bargaining across different types of conflict.

Second, I test whether relative reinforcements after a ceasefire is signed have a pacifying or destabilizing effect on ceasefire survival (hypothesis 2). Using a sample of ceasefires established after interstate wars of the 20th century, I find that conflicts of a high level of violence are less likely to recur when shifts toward imbalance happen. Also, when reinforcement creates parity, peace lasts longer. Given that in the model a rejection of A's offer leads to a battle, the lack of relationship between nonviolent conflict and patterns of reinforcements does not undermine my model.

Finally, in chapter 3 I develop and test a software application designed to address a large parameter space problem in simulation models. A large parameter space problem is at the root of the difficulties that analysts experience with verifying and validating simulation models. The program provides communication between a genetic algorithm and a simulation model, treating the numerical output from the simulation model as a fitness function of the genetic algorithm. The program relies on a genetic algorithm to evolve the initialization parameter values for simulation models until the performance of the model matches a certain criterion of fitness by recombining the best solutions from the previous generations to improve performance and by mutating some solutions' elements to avoid being stuck in a flat region of the space. Optimization heuristics like genetic algorithms infer global optima from evaluating a partial parameter space. This means that the genetic automated search program is a more efficient way of analyzing simulation models in a systematic fashion.

6.2 Contributions to the Literature

This project contributes to multiple strands of literature on third-party conflict management. I offer new conclusions with respect to the role of external support in inter- and intra-state war duration by looking at how reinforcement changes power structure of the dyad as opposed to the identity of the recipient of support. In addition, my model suggests that the literature on peace stability after power mediation has overlooked an important factor that shapes ceasefire duration – external support provided after a ceasefire is signed.

6.2.1 Literature on external support and war duration:

The researchers of external support in civil and interstate conflict have measured support as assistance to a specific side in the conflict. While the central questions around the role of external support differ between inter- and intra-state war literatures, the measurements of support in both lines of research are based on the identity of the support's recipient: rebels vs. governments in the civil conflict literature and initiators vs. targets in the international conflict literature. My model yields that the identity of the recipient of support is irrelevant. Instead, it is the direction in which the power structure of the dyad changes as a result of support that determines whether support will prolong or shorten conflict. Specifically, when reinforcement creates greater parity, it extends conflict, and when support creates greater imbalance, it shortens violence. Reinforcement has a bimodal effect in my model, because when aid changes the relative capabilities in the direction of greater imbalance, the difference in benefits between defeating two different types of opponent diminishes and therefore the offer-making side is more likely to give up more of the prize without fighting. In contrast, when reinforcement creates greater parity in the dyad, the benefits from defeating a weaker as opposed to a stronger type of opponent appear greater to the offer-making side and she is more likely to initiate war or risk another battle in an ongoing war. Thus, my model offers a more general argument that can be applied to both inter- and intra-state conflict. A statistical test introduced in section 5.2 against a sample of post-WWII civil wars is consistent with my model.

The literature on the effect of external support on civil war duration has established that support for the rebels shortens the time until rebel victory, and support for the government prolongs it.[2] In addition, intervention on the government's side either increases the time until government's victory (Licht 2011) or has no effect on it (Gent 2008). Given that the specific questions asked in the literature slightly differ from the question that is central to this project (duration until outcome vs. duration), I reestimated the results in Gent (2008) such that the dependent variable is time until the end of conflict. My analysis in section 5.2 shows that

would allow me to test the model's hypothesis 2, i.e. what is the effect of external support on the conflict's duration, instead of conflict's outcome.

In civil war studies, the debate has been on whether support for the rebels as opposed to the government prolongs war or shortens it. In interstate war research, the major question has been what characteristics of third parties predict joining on the side of the target as opposed the initiator.

In terms of the effects of outside support on conflict duration, there have been inconsistent findings in the civil war literature. Regan and colleagues conclude that the aggregated indicators of external support extend the duration of civil wars (Regan 2000, 2002; Regan

[2] See the summary of literature in subsection 4.1.2, specifically, Regan (2000, 2002); Regan and Aydin (2006); Gent (2008) for the results applicable to the post-WWII civil conflicts and Balch-Lindsay, Enterline and Joyce (2008); Licht (2011) for the results applicable to COW-defined civil wars in 1818-2005.

and Aydin 2006). Two studies – Balch-Lindsay, Enterline and Joyce (2008)[3] and Collier, Hoeffler and Soderbom (2004)[4] – have found that dichotomized measures of external support shortens civil war duration. However, the finding in Balch-Lindsay, Enterline and Joyce (2008) is explained by the incomplete interpretation of the effects over time (Licht 2011). After correct interpretation, joining on both government and rebel side extends war duration in the Correlates of War sample. Collier, Hoeffler and Soderbom (2004), on the other hand, derive their conclusion due to the assumption that the effect of external support inflates rather than decays over time (Regan 2010).[5] Therefore, empirical models yield that external support prolongs civil conflict.

Unlike other studies, Gent (2008) accounts for the supply of support (see the discussion further below) and demonstrates that military reinforcement of the government has no effect on the duration until government's victory, yet support to the rebels shortens conflict; Gent argues that this finding supports his game-theoretical exercise, because it conforms to the notion that support is given when the recipient cannot win without it. The author makes the case further that, in terms of reinforcing the government side, this means that the government was losing without external support, and in terms of bolstering rebels, it means that the rebels were close to winning. Gent's (2008) contribution is twofold. First, the author develops a formal model of the supply side of external support. Second, the author discusses his empirical findings about war duraion in terms of a selection effect of reinforcement and the process of fighting, although he does not account for the selection effect empirically.

[3] Balch-Lindsay, Enterline and Joyce (2008) use the Correlates of War (COW) sample of civil wars and military interveners.

[4] Collier, Hoeffler and Soderbom (2004) use Regan's (2000, 2002) sample of civil wars and interveners.

[5] Modeling the effect of intervention as a decaying effect over time has been standard in the literature.

While significant transfers of weapons and equipment are commonplace in interstate wars,[6] there has been a lack of quantitative research on the effect of such transfers on interstate war duration. The effect of multiple states participating in war as opposed to a two-side war duration has been shown to reduce interstate wars (Bennett and Stam 1996). However, participation of multiple states in an interstate war should not be viewed as 100% analogous to sending troops to support one of the sides in a civil war, as some of the time multiple states are original belligerents as opposed to joiners. Since I assume that the underlying logic of conflict is analogous in both intra- and inter-state wars, I expect my model's predictions to apply to both "types" of conflict as long as I account for the characteristics that are more often observed in some conflicts rather than others.[7]

Recall that the previous research has been investigating the impact of identity-based interventions on duration until a certain outcome, therefore the comparison with Balch-Lindsay, Enterline and Joyce (2008); Gent (2008) is not straightforward.[8] Both works have established that rebel-biased intervention shortens the time until opposition victory and civil war terminations in a negotiated settlement. Gent (2008) finds that government-biased intervention has no effect on war duration until government victory and Licht (2011) reestimates the findings in Balch-Lindsay, Enterline and Joyce (2008) to find that government-biased intervention prolongs the duration of war until government victory. To make the outcomes of the

[6] E.g., consider the US supply of weapons and equipment to Israel during the Yom Kippur War in 1973, or the USSR's weapons supply to the Arab states on the opposing side.

[7] See subsection 1.2 for a more detailed explanation of why I consider the logic of intra- and inter-state wars to be analogous.

[8] These works represent the latest analyses for the most commonly data used in this research: Regan's data in Gent's study and COW Inter-State War data in the study by Balch-Lindsay and colleagues.

Given that most wars that experience intervention receive multiple instances of intervention, it is important to separate the net effects of intervention on each side by introducing a multiplicative interaction of the two intervention types. To understand the effect of each type I use linearly combined coefficients. When estimating duration until the end of war as opposed to one side's victory over another, both types of intervention have a decreasing effect on the hazard of failure. The linearly combined coefficient for government-biased intervention is -1.2, which is statistically significant at the 0.5 statistical significance level. The linearly combined coefficient for the rebels-biased intervention is -0.73, and it is not statistically significant. Similarly, when the intervention types are recoded as support of the stronger or weaker side, both linearly combined effects indicate that both types of intervention prolong war. The effects of control variables are consistent with the published studies: wars with higher death tolls are more likely to end sooner, while ethnic/religious wars take longer to end. Also, the Cold War indicator suggests that during the Cold War civil conflicts were more likely to end sooner. A more difficult terrain is associated with longer wars (although not significant in model 1 of Table 5.1), and political regime has no effect on war duration.

Overall, the results suggest a few takeaways. First, my taking into account of the selection stage that leads to third-party reinforcement generates substantially different results with respect to whether reinforcement prolongs war or shortens it. Specifically, the effects of identity-based and power-based measures of intervention appear to be analogous when selection is not taken into account, yet these measures generate different results when selection is accounted for. Second, Gent (2008) acknowledges that sometimes rebels dominate the government but states the percentage of such cases is insignificant. I demonstrate that recoding support indicators based on power structure in the dyad generates results quite different from identity-based indicators in the full information duration model. These results are consistent with the general bargaining model of conflict management, which is suggestive of civil wars exhibiting analogous bargaining dynamics with other types of conflicts.

6.2.2 Literature on power mediation and peace duration:

With respect to the effect of power mediation on peace duration, my model disrupts the existing convention by allowing for provisions of external support after ceasefire has been signed. While such reinforcements are commonly observed, quantitative research on mediation and peace stability has ignored the provisions of reinforcement after ceasefire as a variable that may affect peace duration. My model suggests

Hypothesis 2: Net reinforcements toward greater imbalance prolong peace and net reinforcements toward greater parity destabilize peace.

Second, while power mediation "delays the inevitable" by creating unstable ceasefires, the survival of imposed ceasefires may be shaped by external support. Power mediation operates through inflating the costs of fighting. When provided during the conflict of interests stage, the sides observe great costs of fighting and derive that fighting is prohibitively costly and return to the peaceful stage for as long as pressure lasts.[9] Once the costs deflate to their previous amount, the sides will again find themselves dissatisfied with the status quo if no redistribution of the disputed good takes place. The reason why redistribution may take

[9] This means that one's assumption about the duration of mediation's effect is critical to the conclusions one could derive about the effectiveness of power mediation; see section 6.3 for the discussion of possible approaches to this problem.

place in my model is the delayed materialization of support provided during the conflict of interests stage.

The effect of power mediation lasts as long as the mediator is willing to generate the barriers to fighting, whether it is threats of sanctions, potential of intervention, physical barriers to fighting, or financial incentives not to fight. In the extended model, the non-waning pressure that generates longer peace spells is the result of repeated instances of provision of power mediation.[10][11]

6.3 Planned Extensions and Changes

6.3.1 The agent-based model of conflict management:

A few immediate changes are necessary for the extended agent-based model. First, I plan to incorporate the costs of fighting into third parties' calculation of an optimal size of investment. Right now, third parties only evaluate how much the desired policy outcome diverges from the current status quo, what amount it would take to help the recipient win to achieve the desired outcome, and whether the third party controls enough resources to afford such reinforcement (see equation 4.1). When power mediators invest, they inflate the costs of fighting. In the current version of the model, support given in the same time period when power mediation is applied has little to no effect until power mediation's effect wanes. Another approach to modeling this dynamic would be to allow potential supporters

[10] In reality, the same mediator could also exert pressure repeatedly and extend peace. Although my model does not explicitly prohibit the same mediator from exerting repeated pressure, it is the order in which third parties decide whether to mediate that makes the same agent unlikely to be randomly chosen to mediate repeatedly.

[11] My model is consistent with current research on third parties prolonging peace through creating additional barriers to fighting or incentives to avoid fighting. For an example of barriers to fighting see work by Fortna (2004, 2008) on peacekeeping. For an example of third parties providing financial incentives to avoid fighting, see Arena and Pechenkina (N.d.) on US aid to Israel and Egypt.

to observe the costs of fighting which would mean that some supporters would be deterred by power mediation.

Second, in the current version of the extended model, the cost of power mediation is proportional to the dyad's total capability and is inversely proportional to the amount of resources controlled by the third party in the triad (see equation 4.3). In the current setup, there is no difference between the concept of the cost of mediation and the ability to exert pressure. Therefore, power mediators sometimes pay great cost to inflate the costs of fighting for the belligerents. I believe a more realistic approach would be to use the current cost of power mediation as a screening device (to screen out the third parties that are not powerful enough to pressure the belligerents into settlement) and then subtract a smaller amount of capability from a third party as the actual cost of one's pressure for peace shouldn't be as high as the potential cost of following through on one's threats.

Third, I would like to introduce random shocks to capabilities in the system to simulate the discovery of new weapons or natural resources. In addition, this option would allow me to compare the influence of third parties and the influence of random shocks on conflict recurrence.

Finally, I plan to add an option of feeding in the initialization parameter values via a graphical user interface to facilitate the use of the model by other analysts and especially to facilitate the use of the model as a tool for generating intuition about real world events. For instance, the current conflict in Syria involves multiple fighting dyads (Syrian government vs. multiple rebel groups) and multiple third parties providing reinforcements to different actors of the conflict. One could use my agent-based model for generating intuition about implications of the number of third parties involved or the divergence of third parties' preferences or the amounts of support third parties provide.

6.3.2 The automated genetic search software program:

The introduced version of the *SimGA* software program yields multiple solutions that satisfy the sought fitness conditions. In the context of simulation models, this means that the genetic search program finds multiple strings of initialization parameter values that generate the emergent property of the simulation model that satisfies the criterion specified by an analyst. Given that simulation models are inherently stochastic an analyst should run the genetic search program multiple times (in the discussed applications in chapters 3 and 4 I have run the program 500 times) to ensure that a representative sample of solutions is found. In the best case scenario, all solutions share common trends: the same model parameters take on high/low values consistently across all solutions. However, when solutions differ among themselves, an analyst would need to employ a pattern recognition algorithm to determine how the model parameters interact to generate the emergent properties of the model. At the moment, the *SimGA* software does not include a pattern recognition module that would allow for an automated assessment of patterns across multiple solutions. Including such a module would improve the ease of interpretation of *SimGA*'s output. The genetic search software could benefit from an addition of, among others, a neural network or a naive Bayes classifier module.[12]

[12] In addition, to speed up convergence it is (Yen et al. 1998).

The current version of the code that implements the reading and interpretation of the simulation model's output is only able to parse through n rows of data recorded to a text file and calculate a summary statistic of a specific variable. In other words, the application *SimGA* currently only interprets cross-sectional numerical output from simulation models. Some models require time series cross-sectional (TSCS) numerical output to reflect the evolution of interactions (e.g., the extended model would be much easier to analyze if *SimGA* could read in the TSCS data format). I would like to rewrite the portion of code for the genetic search application to enable it to manipulate multiple files to increase the flexibility of the application.

While I have demonstrated that the program is effective at reaching sought fitness, the current termination criterion is for one candidate solution to match the desired target value of the response variable. This is a commonplace practice for termination conditions. A more reliable termination condition would be to end the execution if a threshold percentage of candidate solutions meets the sought criterion.

In the current discussion of the program's performance, I do not test for the reliability of the genetic algorithm embedded in the automated search application. Reliability refers to the algorithm's capacity to generate the same result when the algorithm's parameters are varied (Reed N.d.). In other words, to evaluate the reliability of the algorithm employed in *SimGA*, I would have to vary the parameters P, ρ_c, ρ_m, explained in Table 3.3.

Finally, in my future work, I would like to incorporate a multi-objective genetic algorithm into the automated search software program. The use of multi-objective optimization has already revolutionized the way water sharing/supply and airline scheduling are executed

(Reed N.d.). Such a change would allow an analyst to optimize n criteria of the model at a time, thus allowing for a more comprehensive analysis of model dynamics.

6.3.3 Empirical analysis:

Unlike in my analysis of hypothesis 1, examination of hypothesis 2 against a sample of interstate war ceasefires does not take into account the selection causes of the US providing external reinforcement to a particular recipient. I will need to generate predicted hazards of the provision of the US military aid and then include that in the probit models shown in Table 5.3.

I would also want to reestimate my conclusions about testing hypothesis 1 against the data in Cunningham, Gleditsch and Salehyan (2009), as those data provide a different coding scheme for rebel capabilities.

Finally, I would like to merge external support data into the sample introduced in **?** that combines civil and international war data.

Bibliography

Arena, Philip and Anna O. Pechenkina. N.d. “External Subsidies and Lasting Peace.” Presented at the Midwest Political Science Association, Chicago, IL, April, 2011.

Axelrod, Robert. 2003. “Advancing the Art of Simulation in the Social Sciences.” *Japanese Journal for Management Information System* 12(3).

Axelrod, Robert. 2006. Agent-Based Modeling as a Bridge between Disciplines. In *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, ed. Leigh Tesfatsion and Kenneth L. Judd. Handbooks in Economics Series, North-Holland/Elsevier, Amsterdam.

Bäck, Thomas, David B. Fogel and Zbigniew Michalewicz, eds. 2000. *Handbook of Evolutionary Computation*. IOP Publishing Ltd. and Oxford University Press.

Balch-Lindsay, Dylan, Andrew J. Enterline and Kyle A. Joyce. 2008. “Third-Party Intervention and the Civil War Process.” *Journal of Peace Research* 45(3):345–363.

Beardsley, Kyle C. 2008. “Agreement without Peace? International Mediation Agreement without Peace? International Mediation and Time Inconsistency Problems.” *American Journal of Political Science* 52(4):723–740.

Beardsley, Kyle C. 2011. *The Mediation Dilemma*. Cornell University Press.

Beardsley, Kyle C., David M. Quinn, Bidisha Biswas and Jonathan Wilkenfeld. 2006. “Mediation Style and Crisis Outcomes.” *Journal of Conflict Resolution* 50(1):58–86.

Beck, Nathaniel, Jonathan Katz and Richard Tucker. 1998. “Taking Time Seriously: Time-Series-Cross-Section Analysis with a Binary Dependent Variable.” *American Journal of Political Science* 42(4):1260–1288.

Bennett, D. Scott. 2006. “Towards a Continuous Specification of the Democracy-Autocracy Connection.” *International Studies Quarterly* 50(2):313–338.

Bennett, D Scott and Allan Stam. 1996. “The Duration of Interstate Wars, 1816-1985.” *American Political Science Review* 90(2):239–257.

Benoit, Kenneth. 2001. “Simulation Methodologies for Political Scientists.” *Political Methodologist* 10(1):12–16.

Bercovitch, Jacob and Scott Sigmund Gartner. 2006. “Is There Method in the Madness of Mediation? Some Lessons for Mediators from Quantitative Studies of Mediation.” *International Interactions* (32):329–354.

Bercovitch, Jacob and Scott Sigmund Gartner. 2009. Is there method in the madness of mediation? Some lessons for mediators from quantitative studies of mediation. In *Interational Conflict Mediaton. New Approaches and Findings*, ed. Jacob Bercovitch and Scott Sigmund Gartner. Routledge.

Bianchi, Carlo, Pasquale Cirillo, Mauro Gallegati and Pietro A. Vagliasindi. 2007. "Validating and Calibrating Agent-Based Models: A Case Study." *Computational Economics* 30(3):245–364.

Boehmke, Frederick J., Daniel S. Morey and Megan Shannon. 2006. "Selection Bias and Continuous-Time Duration Models: Consequences and a Proposed Solution." *American Journal of Political Science* 50(1):192–207.

Box-Steffensmeier, Janet M. and Christopher J.W. Zorn. 2001. "Duration Models and Proportional Hazards in Political Science." *American Journal of Political Science* 45:951–67.

Box-Steffensmeier, Janet M., Dan Reiter and Christopher J.W. Zorn. 2003. "Nonproportional Hazards and Event History Analysis in International Relations." *Journal of Conflict Resolution* 47(1):33–53.

Brambor, Thomas, William Roberts Clark and Matt Golder. 2006. "Understanding Interaction Models: Improving Empirical Analyses." *Political Analysis* 14:63–82.

Brenner, Thomas and Claudia Werker. 2007. "A Taxonomy of Inference in Simulation Models." *Computational Economics* 30(3):227–244.

Carley, Kathleen M. N.d. "Validating Computational Models." Working Paper, Carnegie Mellon University, September 1996.

Carter, David and Curtis Signorino. 2010. "Back to the Future: Modeling Time Dependence in Binary Data." *Political Analysis* 18(3):271–292.

Cederman, Lars-Erik. 2001. "Agent-Based Modeling in Political Science." *Political Methodologist* 10(1):16–23.

Coello Coello, Carlos A., Gary B. Lamont and David A. Van Veldhuizen. 2007. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Second ed. Springer.

Collier, Paul, Anka Hoeffler and Mans Soderbom. 2004. "On the Duration of Civil War." *Journal of Peace Research* 41(3):253–273.

Cunningham, David E. and Douglas Lemke. 2013. "Combining Civil and Interstate Wars." *International Organization* 3(67):609–627.

Cunningham, David E., Kristian Skrede Gleditsch and Idean Salehyan. 2009. "It Takes Two: A Dyadic Analysis of Civil War Duration and Outcome." *Journal of Conflict Resolution* 53(4):570–597.

DeRouen, Karl, Jacob Bercovitch and Peter Wallensteen. N.d. "Mediation and the Recurrence of Civil War." Presented at the Annual Meeting of International Studies Association, April 1-5, 2012, San Diego, CA.

DeRouen, Karl R. and David Sobek. 2004. "The Dynamics of Civil War Duration and Outcome." *Journal of Peace Research* 41(3):303–320.

Downs, George W., David M. Rocke and Peter N. Barsoom. 1996. "Is the Good News about Compliance Good News about Cooperation?" *International Organization* 50(3):379–406.

Entous, Adam. June 15, 2013. "Behind Obama's About-Face on Syria." *Wall Street Journal* .

Fagiolo, Giorgio, Alessio Moneta and Paul Windrum. 2007. "A Critical Guide to Empirical Validation of Agent-Based Models in Economics: Methodologies, Procedures, and Open Problems." *Computational Economics* 30(3):195–226.

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

Fearon, James D. 1998. "Bargaining, Enforcement, and International Cooperation." *International Organization* 52(2):269–305.

Fearon, James and David Laitin. 2003. "Ethnicity, Insurgency and Civil War." *American Political Science Review* 97(1):75–90.

Fey, Mark and Kristopher Ramsay. 2010. "When is Shuttle Diplomacy Worth the Commute? Information Sharing Through Mediation." *World Politics* 62(4):529–560.

Filson, Darren and Suzanne Werner. 2002. "A Bargaining Model of War and Peace: Anticipating the Onset, Duration, and Outcome of War." *American Journal of Political Science* 46(4):819–37.

Fisher, R.J. 1995. Pacific, Impartial Third Party Intervention in International Conflict: a Review and Analysis. In *Beyond Confrontation: Learning Conflict Resolution in the Post Cold War Era*, ed. John A. Vasquez, Sanford Jaffe, James Turner Johnson and Linda Stamato. University of Michigan Press pp. 39–59.

Fortna, Virginia Page. 2002. "The Cease-Fires Data Set." Time Varying Covariates Version. **URL:** *http://www.columbia.edu/ vpf4/research.htm*

Fortna, Virginia Page. 2003. "Scraps of Paper? Agreements and the Durability of Peace." *International Organization* 57:337–372.

Fortna, Virginia Page. 2004. *Peace Time: Cease-Fire Agreements and the Durability of Peace*. Princeton, NJ: Princeton University Press.

Fortna, Virginia Page. 2008. *Does Peacekeeping Work?: Shaping Belligerents Choices after Civil Wars*. Princeton, NJ: Princeton University Press.

Galán, José Manuel, Luis R. Izquierdo, Segismundo S. Izquierdo, José Ignacio Santos, Ricardo del Olmo, Adolfo López-Paredes and Bruce Edmonds. 2009. "Errors and Artefacts in Agent-Based Modeling." *Journal of Artificial Societies and Social Simulation* 12(1):1. **URL:** *http://jasss.soc.surrey.ac.uk/12/1/1.html*

Gartzke, Erik. 1999. "War Is in the Error Term." *International Organization* 53(3):567–87.

Gartzke, Erik. 2007. "The Capitalist Peace." *American Journal of Political Science* 51(1):166–191.

Gent, Stephen. 2008. "Going in When It Counts: Military Intervention and the Outcome of Civil Conflicts." *International Studies Quarterly* 52(4):713–735.

Ghosn, Faten, Glenn Palmer and Stuart Bremer. 2004. "The MID3 Data Set, 1993–2001: Procedures, Coding Rules, and Description." *Conflict Management and Peace Science* 21:133–154.

Givens, Geof H. and Jennifer A. Hoeting. 2005. *Computational Statistics*. Wiley Series in Probability and Statistics: John Wiley & Sons, Inc., Hoboken, New Jersey.

Greig, J. Michael and Patrick M. Regan. 2008. "When Do They Say Yes? An Analysis of the Willingness to Offer and Accept Mediation in Civil Wars." *International Studies Quarterly* (58):759781.

Headlee, Celeste. June 19, 2013. "To Arm, or Not to Arm the Syrian Rebels? An Interview with Shadi Hamid." *National Public Radio* . **URL:** *www.npr.org/templates/story/story.php?storyId=193495486*

Hirshleifer, Jack. 2000. "The Macrotechnology of Conflict." *Journal of Conflict Resolution* 44(6):773–792.

Horner, J., M. Morelli and F. Squintani. N.d. "Mediation and Peace." Working Paper.

How to Use Parameters and Parameter files. RePast Documentation. N.d. Web resource. **URL:** *http://repast.sourceforge.net/repast3/how-to/params.html*

Huth, Paul K. 1998. "Major Powers Intervention in International Crises, 1918-1988." *Journal of Conflict Resolution* 42(6):744–770.

Joyce, Kyle A. 2008. "The Agent-Based Model of War Expansion." PhD thesis The Pennsylvania State University.

Joyce, Kyle A., Faten Ghosn and Reşat Bayer. 2013. "When and Whom to Join: The Expansion of Ongoing Violent Interstate Conflict." *British Journal of Political Science* February:1–34.

Kydd, Andrew. 2003. "Which Side Are You On? Bias, Credibility, Which Side Are You On? Bias, Credibility, and Mediation." *American Journal of Political Science* 47(4):597–611.

Licht, Amanda. 2011. "Change Comes with Time: Substantive Interpretation of Nonproportional Hazards in Event History Analysis." *Political Analysis* 19:227–243.

Lo, Nigel, Barry Hashimoto and Dan Reiter. 2008. "Ensuring Peace: Foreign-Imposed Regime Change and Postwar Peace Duration, 1914 - 2001." *International Organization* 62(Fall):717–736.

Luttwak, Edward N. 1999. "Give War a Chance." *Foreign Affairs* 4(78):36–44.

Marks, Robert Ernest. 2007. "Validating Simulation Models: A General Framework and Four Applied Examples." *Computational Economics* 30:265–290.

Meirowitz, Adam, Massimo Morelli, Kristopher W. Ramsay and Francesco Squintani. N.d. "Mediation and Strategic Militarization." Working Paper.

Melin, Molly M. 2011. "The Impact of State Relationships on If, When, and How Conflict Management Occurs." *International Studies Quarterly* (55):691715.

Melin, Molly M. and Michael T. Koch. 2010. "Jumping into the Fray: Alliances, Power, Institutions, and the Timing of Conflict Expansion." *International Interactions* 1(36):1–27.

Midgley, David F., Robert E. Marks and Dinesh Kunchamwar. 2007. "The building and assurance of agent-based models: An example and challenge to the field." *Journal of Business Research* 60:884–893.

Miller, John H. 1998. "Active Nonlinear Tests (ANTs) of Complex Simulation Models." *Management Science* 44(6):820–830.

Miller, John H. and Scott E. Page. 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press.

Mitchell, Melanie. 1997. *An Introduction to Genetic Algorithms*. MIT Press.

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

Powell, Robert. 1999. *In the Shadow of Power: States and Strategies in International Politics*. Princeton University Press.

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

Powell, Robert. 2006. "War as a Commitment Problem." *International Organization* 60(1):169–203.

Primo, David and Kevin Clarke. 2007. "Modernizing Political Science: A Model-Based Approach." *Perspectives on Politics* 5(4):741–753.

Quinn, David, Jonathan Wilkenfeld, Pelin Eralp, Victor Asal and Theodore McLauchlin. 2013. "Crisis Managers but Not Conflict Resolvers: Mediating Ethnic Intrastate Conflict in Africa." *Conflict Management and Peace Science* 4(30):387–406.

Reed, Patrick M. N.d. "Evolutionary Algorithms for Multiobjective Optimization." 2013 Course Lectures. Penn State University.

Reed, William. 2003. "Information, Power, and War." *American Political Science Review* 97(4).

Reed, William, David H. Clark, Timothy Nordstrom and Wonjae Hwang. 2008. "War, Power, and Baargaining." *The Journal of Politics* 70(4):1203–1216.

Regan, Patrick and Aysegul Aydin. 2006. "Diplomacy and Other Forms of Intervention." *Journal of Conflict Resolution* 50(5):736–756.

Regan, Patrick M. 2000. *Civil Wars and Foreign Powers: Interventions and Intrastate Conflict*. University of Michigan Press.

Regan, Patrick M. 2002. "Third-Party Interventions and the Duration of Intrastate Conflicts." *Journal of Conflict Resolution* 46(1):55–73.

Regan, Patrick M. 2009. "Civil War and Territory?: Drawing Linkages Between Interstate and Intrastate War." *International Interactions* 35(3):321–329.

Regan, Patrick M. 2010. "Interventions into Civil Wars: A Retrospective Survey with Prospective Ideas." *Civil Wars* 12(4):456–476.

Regan, Patrick M., Richard Frank and Aysegul Aydin. 2008. "Diplomatic Interventions and Civil Wars: A New Dataset." *Journal of Peace Research* .

Richiardi, Matteo, Roberto Leombruni, Nicole Saam and Michele Sonnessa. 2006. "A Common Protocol for Agent-Based Social Simulation." *Journal of Artificial Societies and Social Simulation* 9(1).
URL: *http://jasss.soc.surrey.ac.uk/9/1/15.html*

Sambanis, Nicholas. 2004. "What Is Civil War?: Conceptual and Empirical Complexities of an Operational Definition." *Journal of Conflict Resolution* 48(6):814–858.

Sarkees, Meredith Reid. 2000. "The Correlates of War Data on War: An Update to 1997." *Conflict Management and Peace Science* 18(1):123–144.

Sarkees, Meredith Reid and Frank Wayman. 2010. *Resort to War: 1816 - 2007*. CQ Press.

Schelling, Thomas C. 1978. *Micromotives and Macrobehavior*. W. W. Norton & Company, Inc.

Schrodt, Philip A. and Deborah J. Gerner. 2004. "An Event Data Analysis of Third-Party Mediation in the Middle East and Balkans." *Journal of Conflict Resolution* 48(3):310–330.

Singer, J. David, Stuart Bremer and John Stuckey. 1972. Capability Distribution, Uncertainty, and Major Power War, 1820-1965. In *Peace, War, and Numbers*, ed. Bruce Russett. Beverly Hills: Sage pp. 19–48.

Slantchev, Branislav L. 2003. “The Principle of Convergence in Wartime Negotiations.” *American Political Science Review* 97(4):621–32.

Smith, Alastair and Allan Stam. 2004. “Bargaining and the Nature of War.” *Journal of Conflict Resolution* 48(6):783–813.

Terris, Lesley G. and Zeev Maoz. 2005. “Rational Mediation: a Theory and a Test.” *Journal of Peace Research* 5(42):563–583.

Von Stein, Jana. 2005. “Do Treaties Constrain or Screen? Selection Bias and Treaty Compliance.” *American Political Science Review* 99(04):611–622.

Wagner, Harrison R. 2000. “Bargaining and War.” *American Journal of Political Science* 44(3):469–484.

Walter, Barbara. 1997. “The Critical Barrier to Civil War Settlement.” *International Organization.* 51(3).

Werner, Suzanne and Amy Yuen. 2005. “Making and Keeping Peace.” *International Organization* 59(Spring):261–292.

Wilkenfeld, Jonathan, Kathleen Young, Victor Asal and David M. Quinn. 2003. “Mediating International Crises. Cross-National and Experimental Perspectives.” *Journal of Conflict Resolution* 47(3):279–301.

Yen, John, J.C. Liao, B. Lee and D. Randolph. 1998. “A hybrid approach to modeling metabolic systems using a genetic algorithm and simplex method.” *IEEE Transactions on System, Man, and Cybernetics* 2(28):173–191.

Zakaria, Fareed. June 20, 2013. “Obama’s Syria policy is full of contradictions.” *Washington Post* .

Vita

Anna O. Pechenkina

Education

The Pennsylvania State University University Park, PA 2007–Present
Ph.D. in Political Science, expected in July 2013
Areas of Specialization: International Relations and Methodology

The University of Mississippi Oxford, MS 2003–2006
B.A. in Political Science, *Summa cum laude, Phi Beta Kappa*

Awards and Grants

Pre-Doctoral Teaching Fellowship, Program in Empirical Intl Relations	FA 2011
Summer Graduate Research Grant	SU 2011
Research and Graduate Studies Office Dissertation Release Grant	SP 2011
Jesse M. MacKnight Memorial Graduate Scholarship	SP 2009
Pre-Doctoral Fellowship, Quantitative Social Science Initiative	2008–2009
Miller-LaVigne Graduate Fellowship	2007–2008
Howerton-Fortenberry Award and Governor's Award	SP 2006
Undergraduate Tuition and International Student Scholarships	2004–2006

Research Experience

Doctoral Research The Pennsylvania State University 2010–Present
Thesis Advisor: D. Scott Bennett

This research involved (i) the development of an agent-based model of conflict management (freeware), (ii) the development of a software program to perform an automated genetic search of parameter spaces in simulation models (freeware), and (iii) statistical analysis of the model's propositions against empirical data.

Graduate Research The Pennsylvania State University
Research Advisor: D. Scott Bennett 2009–2013

Contributed to the Insurgency project by programming extensions to existing computer simulations and analyzing the simulation output.

Research Advisor: Glenn Palmer 2007–2008

Contributed to the MID4/Tabari project by identifying the population of conflicts reported in the media for training the automated coding software.

Teaching Experience

Instructor The Pennsylvania State University FA 2011

Developed and taught PL SC 418 "International Relations Theory"