

International Negotiations in the Shadow of Elections

Justin Melnick* and Alastair Smith†

New York University

October 27, 2022

Abstract

Within a formal model of international agreements in the shadow of renegotiations and domestic competition, we highlight three important ways elections shape international agreements. Elections determine who will be in control of policy in the future, which affects how leaders bargain today. Elections also determine the deals policymakers will agree to. Finally, proposers have the opportunity to shape the contours of domestic political competition with what is offered in pre-electoral bargaining. We identify that several canonical results in the literature – like the Schelling conjecture or the idea that hawkish leaders have an innate bargaining advantage over dovish leaders – only hold under certain restrictions on how voters evaluate their leaders. In contrast, we show paradoxically that when voters are prospective, electoral incentives shade the ability for domestic leaders to negotiate better deals for their publics. Counterintuitively, this leads to hawks agreeing to more conciliatory agreements than doves.

Word Count: 10,690

*Email: melnickj@nyu.edu. Phone: 973- 727-6006. Fax: 212-995-4184

†Email: Alastair.Smith@nyu.edu. Phone: 212-998-9678. Fax: 212-995-4184

1 Introduction

International agreements are often negotiated by leaders who are accountable to a domestic audience. We present a formal model of bargaining that considers how the negotiation and renegotiation of international agreements is affected by the shadow of domestic political competition. We examine intertemporal dynamics of optimal policy proposals before and after a domestic election. Many factors influence international negotiations. Our analysis highlights that agreements depend upon the status quo distribution of policy, the likely outcome of the election, and the context in which voters evaluate their leaders in elections.

The preferences of individual leaders matter in shaping foreign policy outcomes. Consider the case of the denuclearization talks between the United States and Iran that led to the signing of the Joint Comprehensive Plan of Action (JCPOA) in 2015. U.S. President Barack Obama hailed the deal as “the strongest non-proliferation agreement ever negotiated,” claiming that “we have achieved a detailed arrangement that permanently prohibits Iran from obtaining a nuclear weapon.”¹ However, on the campaign trail, then-candidate Donald Trump lambasted the terms of the Iran deal and Tehran’s noncompliance, calling the agreement the “highest level of incompetence” and alleging that “we watched them ignore its terms, even before the ink was dry.”² A key theme of Trump’s campaign was his willingness to renegotiate or withdraw from cooperative international agreements: the United States withdrew from the Iran deal in May 2018, with Trump vowing to develop a more comprehensive solution.

We argue that elections interlock with the process of negotiating international agreements in three key ways. First, elections determine who will be in control of foreign policy in the future, which affects what leaders will agree to today. The electoral process invariably injects uncertainty into the stability of international negotiations as countries’ representatives at the bargaining table, and subsequent policy preferences, change. An important consequence of this dynamic is that renegotiation is only credible if a leader inherits terms of an agreement from which he or she would be willing to walk away. Trump’s more “hawkish” inclinations versus Obama’s “dovish” approach to international dealmaking meant that the scope for bargaining with the United States under Trump’s leadership narrowed (and subsequently vanished with American withdrawal from the agreement).

A second key factor is that electoral incentives affect the types of deals that policymakers are willing to conclude prior to elections. What leaders will agree to depends on how voters assess their performance and how much they care about remaining in office. Finally, because of these domestic pressures induced by elections, proposers have the opportunity to shape the contours of domestic political competition with what is offered in pre-electoral bargaining. In this way, the concessions outlined in international

¹<https://obamawhitehouse.archives.gov/the-press-office/2015/08/05/remarks-president-iran-nuclear-deal>

²<https://www.nytimes.com/2016/04/28/us/politics/transcript-trump-foreign-policy.html>

agreements directly affect how voters assess their leaders: the electoral outcome is *endogenous* to terms of the agreement. In the 2021 presidential elections, Iranians elected Ebrahim Raisi, a more hawkish leader who ran on insisting greater sanctions relief in exchange for reviving the nuclear agreement. The negotiations and subsequent American withdrawal from the JCPOA emboldened more conservative and hardline forces to succeed electorally in Iran, which in turn affects future bargaining.

We consider a two-period model in which a unitary foreign power F offers to strike an agreement with the leader of country D in exchange for policy concessions. Between periods, there is an election in nation D , in which voters select a dovish leader L or a hawkish leader R . Our primary focus is to illustrate how electoral incentives shape bargaining in the first period. As such, we consider cases in which the incumbent in nation D is a dove or a hawk under different electoral contexts. We show that a foreign power can generally but not always extract greater concessions from doves rather than hawks. However, hawks have better leverage to renegotiate an agreement after an election because they can credibly threaten to walk away from the table when a dove could not. Paradoxically, as we will show, if voters evaluate candidates prospectively, then this “advantage” can lead hawks to reject deals in the short term that are mutually advantageous to both sides, and in extreme cases, agree to worse deals than doves.

Our model casts negotiations of international arrangements against the backdrop of domestic politics. More generally, we build on the literature examining the role of domestic politics in facilitating or inhibiting international negotiations. Early work in this domain emphasized the “two-level game” structure (Putnam 1988) and sought to understand how domestic actors served as ratification constraints in bargaining at the global level. The so-called “Schelling conjecture” (Schelling 1960; Milner 1997) claimed that domestic ratification constraints provide negotiators with greater leverage in international bargaining (Tarar 2001).

Formal ratification of international agreements is rare. Instead, we argue that concerns about international negotiations are brought up as issues to the domestic electorate, to whom a leader must cater in order to remain in office. We present a model of elections and highlight how different electoral contexts affect which agreements are forged and the domestic implications of these agreements (cf. Battaglini and Harstad 2020; Buisseret and Bernhardt 2018; Milner and Rosendorff 1997; Smith and Hayes 1997). In similar models, voters or other endogenous domestic constraints behave like veto players within the two-level game structure (Iida 1993; 1996; Mo 1994; Chapman, Urpelainen and Wolford 2013; Böhmelt 2019). We depart from the traditional ratification literature by considering how domestic leaders value the relative tradeoff between conceding foreign policy to forge an agreement and remaining in office.

Our focus on the incentives of leaders also contributes to the discussion of how renegotiation of international outcomes relies on preferences of individual actors. The threat of leader turnover affects

the calculus of the foreign power in its original offer. Depending on electoral prospects, a foreign power may choose to withhold international cooperation (Smith 2009) to exploit all gains from bargaining after the election. Other models emphasizing political survival and turnover cast this as a commitment problem (Wolford 2012), but we internalize the risk of renegotiation into the players' strategic calculus directly. In a similar fashion, Brown and Urpelainen (2015) argue that treaty negotiators can shape political outcomes by strategically mobilizing domestic interest groups. Leaders thus must weigh how accepting or rejecting an agreement may affect their electoral prospects.

Our model considers different electoral contexts (Ferejohn 1986). We examine cases of retrospective voting, in which the electorate assesses an incumbent's past performance (Lanoue 1994; Healy and Malhotra 2013), as well as prospective voting, whereby voters factor in how an agreement may be renegotiated depending upon who wins the election (Reed and Cho 1998).

We also contribute to the literature on how hawkish and dovish leaders impact international bargaining. The mere presence of a domestic opposition has been theorized to be a means of deterring conflict (Schultz 1998; Ramsay 2004), but there is also a large literature examining specific leaders' propensities for conflict and cooperation (Horowitz, Stam and Ellis 2015). Several studies have argued that hawks, particularly moderate hawks, are the most likely actors to bring about sustained cooperation and receive greater electoral support for doing so (Cowen and Sutter 1998; Cukierman and Tommasi 1998; Schultz 2005). The origins of these results follow from a hawk's purported policy credibility compared to a dove, and this credibility is particularly likely to result in successful agreement when hawks adopt moderate bargaining positions. By contrast, Clare (2014) argues that doves have bargaining advantages in the face of hawkish oppositions because foreign nations would prefer to negotiate with a dovish leader. Other scholars emphasize reciprocity of cooperation regardless of whether a leader is a hawk or a dove (Colaresi 2004). Empirical evidence has shown that voters prefer hawkish leaders both on credibility and moderation grounds (Mattes and Weeks 2019).

2 Model Description

We consider a potential trade between a domestic nation D and a foreign nation F within a two period setting. While we model F ("it") as a unitary actor, nation D is composed of two leaders/parties, labeled L ("he") and R ("she") and a median voter M ("they"). Between the two periods M decides who they want as their leader for the second period. Hence F potentially deals with a different leader from D in each period. As a modeling assumption, we assume that in the absence of a preexisting deal, F has agenda-setting power and can demand policy concessions from the leader of D in exchange for the execution of a cooperative venture. If an agreement is formed in the first period, then it becomes

the status quo against which any second round renegotiations take place. We suppose that each nation receives a benefit from pursuing cooperative policies, but must bargain over the terms under which such a project would be executed.

The model considers a unidimensional policy $x \in [0, 1]$. The foreign nation's ideal outcome would be $x = 1$, where the domestic nation provides maximal policy concessions in order to implement the cooperative project. Conversely, the domestic nation prefers deals closer to the point $x = 0$, so that it retains as much autonomy over its foreign policy as possible. Although we present this in terms of policy concessions, this could be equally thought of as cost sharing to finance the agreement. Throughout, we will illustrate the core dynamics of the model with quadratic preferences along with a personalistic benefit to pursuing international cooperation. Given terms x , nation F 's payoff from an agreement is

$$\theta_F + v_F(x) = \theta_F - (1 - x)^2$$

where θ_F is the benefit that F receives if an agreement is in place. The payoffs for domestic nation actors $i = L, R, M$ of an agreement with policy x are

$$\theta_i + v_i(x) = \theta_i - x^2$$

where θ_i reflects actor i 's benefit from having a deal in place with nation F . Note that if no agreement is in place, actors do not receive the payoff θ_F or θ_i and simply receive their quadratic utility at the value of the status quo ($x_0 = 0$).

To keep the story simple, we consider the possibility that F can only offer a discrete number of deals to D 's leader. This reflects a simple number of focal possibilities on which countries negotiate with one another. We develop a fuller version of the model in the appendix in which the policy space is continuous. While the discrete choice setting we examine here is not as general as the continuous choice model, it provides the same substantive insights with significantly less complexity.




Figure 1: Policy Space and Possible Agreements

The possible agreements, labeled x_A , x_B , and x_C , are displayed in Figure 1. The foreign power prefers deals in reverse alphabetical order, with x_C being most preferred as it represents the greatest concessions (or cost shares) by D . Players in the domestic nation prefer deals in alphabetical order, with x_A being most preferred. Our critical assumption in the model is that leader L would myopically strike a deal with F on all three possible terms, meaning he would be willing to concede x_A , x_B , or x_C to form

an agreement. Alternatively, R would only myopically be willing to sign away concessions on terms of x_A or x_B .³ This discrepancy in willingness to concede policy implies that R is a more hawkish leader than L . Put differently, L values international cooperation more highly than R , or $\theta_L > \theta_R$. While this assumption is without loss of generality (since we can simply flip party labels), our nomenclature is consistent with the general assumption in the literature (Palmer, London and Regan 2004; Clare 2014) that left-wing parties are more dovish than right-wing parties.

Absent an agreement, the status quo outcome is $x_0 = 0$ and no actor receives their benefit from cooperation. Throughout the paper, we utilize the numerical example in which $x_A = 0$, $x_B = \frac{1}{3}$, and $x_C = \frac{2}{3}$, with $\theta_F = 1$, $\theta_R = \frac{1}{4}$, and $\theta_L = \frac{1}{2}$. We consider both $\theta_M = \frac{1}{4}$ and $\theta_M = \frac{1}{2}$ to explore how outcomes change when the median voter is relatively hawkish or relatively dovish.

In the absence of an existing agreement, we endow F with agenda setting power.⁴ Consistent with standard bargaining models, in the single shot game this would allow F to capture all the bargaining surplus. Once an agreement is established either party might ask to renegotiate the terms. However, no renegotiation proposal will ever be successful unless one of the parties is willing to walk away from the extant agreement and revert to the non-cooperative state. The policy dimension is such that F wants to increase the deal, while D wants to reduce the deal. Since the nations have diametrically opposed preferences with respect to renegotiations, neither will accept the other's renegotiation proposal, unless the proposer can credibly threaten to walk away from the extant agreement. In our setup this occurs when the hawkish leader R is in power in the second period and the existing agreement is x_C . F prefers x_A , x_B and x_C to non-cooperation and so has no ability to credibly renegotiate.

Agreements are not formed in isolation from the domestic political process. When negotiating in the first period, leaders take into account two critical factors: the *impact of today's agreement on future negotiations*; and the *electoral consequences of today's agreement*. International negotiations are thus carried out in the shadow of the polls. To reflect F 's proposal power but inability to renegotiate, we consider the follow sequence of moves.

1. First Period Policy Negotiation:

- (a) F can demand policy concession $x_1 \in \{x_B, x_C\}$ in exchange for the creation of an agreement, or maintain the status quo x_0 .
- (b) D 's leader (L or R) either accepts F 's proposal creating new status quo x_1 , or rejects, maintaining status quo x_0 .

2. Election: M selects either L or R as second period leader.

³In principle, F could offer x_A to either leader but such an action is strictly dominated by offering x_B (see appendix). To simplify exposition, we exclude this possibility from further consideration.

⁴In the interest of brevity, we leave the extension in which D has agenda-setting power to future research.

3. Second Period Policy (Re)negotiation:

- If there is no existing agreement:
 - (a) F can demand policy concession $x_2 \in \{x_B, x_C\}$, or maintain the status quo x_0 .
 - (b) D 's leader (L or R) either accepts x_2 or rejects.
- If there is an existing agreement, then D can propose renegotiation:
 - (a) Domestic leader (L or R) offers a renegotiation, x_2 .
 - (b) F either accepts D 's offer (x_2), retains the existing agreement (x_1), or exits the agreement (x_0).
 - (c) Leader D either remains in the agreement (x_2) or D exits the agreement (x_0).

Players' payoffs are a weighted sum of the payoffs in each period, where δ reflects the relative importance of post-electoral outcomes. One simple interpretation of δ is the time until the election. If the election is far into the future, then δ is relatively small. In contrast, if the election is imminent, then future outcomes are more salient than immediate agreements, which is represented by a large value of δ .

In addition to policy concerns, leaders value office holding, $\Psi > 0$. $\mathbb{I}_{D_2=i}$ is an indicator of whether leader $i = L, R$ is leader of D in the second period. \mathbb{I}_{x_t} is an indicator of whether there is a deal in place in period t . Given the deals x_1 and x_2 for each period, the actors receive payoffs as follows:

$$\begin{aligned}
 U_F(x_1, x_2) &= \mathbb{I}_{x_1} \theta_F + v_F(x_1) + \delta \left(\mathbb{I}_{x_2} \theta_F + v_F(x_2) \right) \\
 U_L(x_1, x_2) &= \mathbb{I}_{x_1} \theta_L + v_L(x_1) + \delta \left(\mathbb{I}_{x_2} \theta_L + v_L(x_2) + \mathbb{I}_{D_2=L} \Psi \right) \\
 U_R(x_1, x_2) &= \mathbb{I}_{x_1} \theta_R + v_R(x_1) + \delta \left(\mathbb{I}_{x_2} \theta_R + v_R(x_2) + \mathbb{I}_{D_2=R} \Psi \right)
 \end{aligned}$$

We pause to emphasize three important features of the model. First, no leader needs to be part of an agreement and the game form provides both F and D the opportunity to opt out of cooperation in each period. Second, leaders are responsible for signing and implementing agreements. So while an agreement will not form unless leader D wants it to, other actors within nation D might prefer that the agreement did not exist. Finally, if an agreement is formed in the first period, then those terms serve as the basis of future negotiations.

2.1 Elections

We examine elections on two dimensions: the *salience* of the international agreement to the voters (σ) and whether the voters evaluate leaders in a retrospective or prospective manner (Ferejohn 1986). We call this element the *electoral context*.

Voters care about more than just foreign policy outcomes (Guisinger 2009; Gadarian 2010). Let u_L and u_R represent the voters' payoff from L 's and R 's negotiations with F . After observing the first period outcome x_1 , voters observe random variables ε_L and ε_R that represent their expectations about the value of L 's and R 's leadership on all other dimensions.

Let $U_M(\text{elect } L) = \sigma u_L + \beta + \varepsilon_L$ be the median voter's payoff from electing L , where β represents any bias in favor of L on all non-policy negotiation issues and σ is the salience of the negotiation within the domestic electoral landscape. Let $U_M(\text{elect } R) = \sigma u_R + \varepsilon_R$ be M 's payoff from electing R . The median voter thus prefers L to R when

$$\varepsilon = \varepsilon_R - \varepsilon_L \leq \beta + \sigma(u_L - u_R)$$

Let $\varepsilon \sim G(\cdot)$, such that the probability that L is elected is $p = G(\beta + \sigma(u_L - u_R))$. We assume G is twice differentiable with full support. Let g be the associated density function, and to keep interpretation simple, we focus on the case where G is uniform such that $p = G(\beta) + g\sigma(u_L - u_R)$. We note that the election probability has two pieces. The first is simply L 's electoral bias $G(\beta)$, and the second is the means through which the policy negotiations affect the election, particularly through the parameter σ . If the election outcome is *exogenous* to the policy negotiations, i.e. $\sigma \rightarrow 0$, then the probability that L is elected to be the second period leader is simply $p = G(\beta)$, which we also denote as \bar{p} .

If $\sigma \neq 0$, the utility garnered by the voters from the agreement depends on how they incorporate it into their electoral calculus. In our running example, we set $\sigma = 1$ and $g = 1$. Retrospective voters evaluate incumbent performance based on a reward-punishment model: if the incumbent has delivered over the course of the term, he is rewarded with reelection. Prospective voters assess electoral candidates based on whether and how they will renegotiate an agreement if they win the election.

Retrospective voters make a simple comparison of what the incumbent leader delivered relative to the status quo. If L is the incumbent, then $u_L = \mathbb{I}_{x_1} \theta_M + v_M(x_1) - v_M(x_0)$, which is the difference in their welfare between a first period agreement and the status quo, and $u_R = 0$. In contrast, if R is the incumbent, then $u_R = \mathbb{I}_{x_1} \theta_M + v_M(x_1) - v_M(x_0)$ and $u_L = 0$. The electorate might be relatively dovish, hence happy with any agreement, or electorate might be relatively hawkish, only willing to make concessions up to x_B .

Conversely, prospective voters base their assessment on what they expect leaders can deliver in the second period. Prospective voters compare their expected welfare under both leaders in the second period given the agreement reached in the first period. Let the notation $\hat{x}_L(x_1)$ represent the second period deal given L is elected and the first round outcome is x_1 . Likewise $\hat{x}_R(x_1)$ corresponds to the second period outcome if R is elected. For prospective voters, $u_L = v_M(\hat{x}_L(x_1))$ and $u_R = v_M(\hat{x}_R(x_1))$.

To put the distinction between retrospective and prospective voters into a clearer perspective, it is

useful to reconsider the Iran nuclear deal. A retrospective voter would compare how much they like the terms of the deal compared to Iran continuing the development of nuclear weapons. If they prefer the deal to the status quo, then they would reward the Democratic candidate in the 2016 election. In contrast, a prospective voter would consider what would be the likely to happen to the Iran deal after the 2016 election. If the Democratic candidate were elected, then the deal would likely continue unchanged. In contrast, Donald Trump, the Republican candidate was very hawkish on the Iran deal. He expressed a willing to terminate the deal, which indeed he did. His willingness to walk away from the extant deal gave him enhanced leverage to demand additional concessions from Iran. Prospective voters who wanted the deal to end, or wanted Iran to agree to harsher terms would have favored Trump, at least to the extent that Iran’s nuclear status influenced their vote (our salience parameter, σ).

To focus on substantively interesting cases, we assume politicians highly value office-holding relative to negotiations such that no politician would voluntarily step down to have another leader negotiate the deal.

2.2 Second Period Outcomes

We categorize the unique subgame perfect equilibrium. The first period deal and election result uniquely determine the policy outcome in the second period.

	No Agreement, $x_1 = x_0$	Agreement, $x_1 \in \{x_A, x_B, x_C\}$		
	x_0	$x_1 = x_A$	$x_1 = x_B$	$x_1 = x_C$
Leader L : $\hat{x}_L(x_1)$	x_C	x_A	x_B	x_C
Leader R : $\hat{x}_R(x_1)$	x_B	x_A	x_B	x_A

Table 1: Second Period Agreements Given First Period Agreement x_1

The results are summarized in Table 1. Recall $\hat{x}_L(x_1)$ and $\hat{x}_R(x_1)$ denote the second period outcomes under L and R ’s leadership given the first round agreement. Pre-existing agreements shape negotiations.

If there is no agreement concluded in the first period, then, in the second period, F can propose an agreement and in equilibrium will leverage its proposal power to maximize policy concessions. That is to say, F obtains the greatest policy concession that D ’s leader is willing to myopically accept. For L , the more dovish leader, this is x_C , while for R , the more hawkish leader, this concession is x_B . This also means that F obtains greater concessions if L is elected.

If an agreement is already in place, then leader D can propose a renegotiation. While D wants greater autonomy over policy concessions (or a smaller cost share), it can only obtain them if he or she can credibly walk away from cooperation. In the discrete case, renegotiation can only occur if the first period outcome is x_C and R wins the election. In this case, R will propose new terms of x_A . Since L is willing to accept all three possible proposals, he cannot credibly walk away from any deal. This

is consistent with other theoretical and empirical findings that hawkish leaders can negotiate better agreements (Buisseret and Bernhardt 2018; Mattes and Weeks 2019).

3 Dovish Incumbent

As summarized in Table 1, the first period deal affects the second period outcome. Thus when negotiating the first period agreement, leaders must consider more than just the immediate policy implications. International negotiations center on three important questions. First, how does signing a deal today affect deals tomorrow? Second, how do electoral incentives affect the domestic incumbent’s willingness to accept or reject deals? Third, how does concluding an agreement affect who will win the election and become the leader in subsequent bargaining?

We start by examining statements of the equilibrium when the domestic leader is relatively dovish. Dovish leader L is myopically willing to agree to greater policy concessions on D ’s behalf in order to secure a deal with F than hawkish leader R ($\theta_L > \theta_R$). In the discrete case, this means that L would accept the deal x_C , while R would not. As we shall see, however, F does not always demand x_C when L is the incumbent, even if this is the maximum concession it could extract. We explore how the introduction of elections creates tensions with L ’s preferences over this policy, and then show how this shapes F ’s optimal offer in the first period.

3.1 Pure Renegotiation Concerns with Doves

When the first period bargaining outcome plays no role in determining the winner of the election, we say that the election is exogenous. Formally, we model this as $\sigma = 0$, and L is elected with probability $p = G(\beta) = \bar{p}$, regardless of the deal concluded. From the perspective of the negotiations, the outcome of the election is akin to a coin flip: electoral incentives do not affect how L accepts or rejects deals at all. In other words, if F were to offer the deal x_C , L would always accept it (i.e., $\theta_L + v_L(x_C) > 0$ if elections are exogenous).

Importantly, the fact that L always accepts x_C does not mean that F will always demand it. Elections create possibilities for renegotiation. If F concludes the agreement x_C with L but R wins the election, then R could renegotiate the terms to x_A . So even though x_C would be the most generous policy concession F could extract, the shadow of potential renegotiations means that it is not always optimal to conclude such a deal.

When the election outcome is exogenous to policy negotiations, F ’s optimal offers to L are shown in Figure 2. We characterize optimal deals as a function of the relative importance of the future, δ , and the electoral bias toward L , β . Unsurprisingly, the largest region of the plot is the purple/striped region, in

which F 's utility-maximizing strategy is to offer x_C . Thus, F often maximally exploits a dovish leader. This deal is myopically acceptable to L , and provides F with large immediate policy guarantees. If L is likely to win the election (β large) and/or the election is far off (δ small), then extracting the maximum sustainable concession from L is F 's best option.

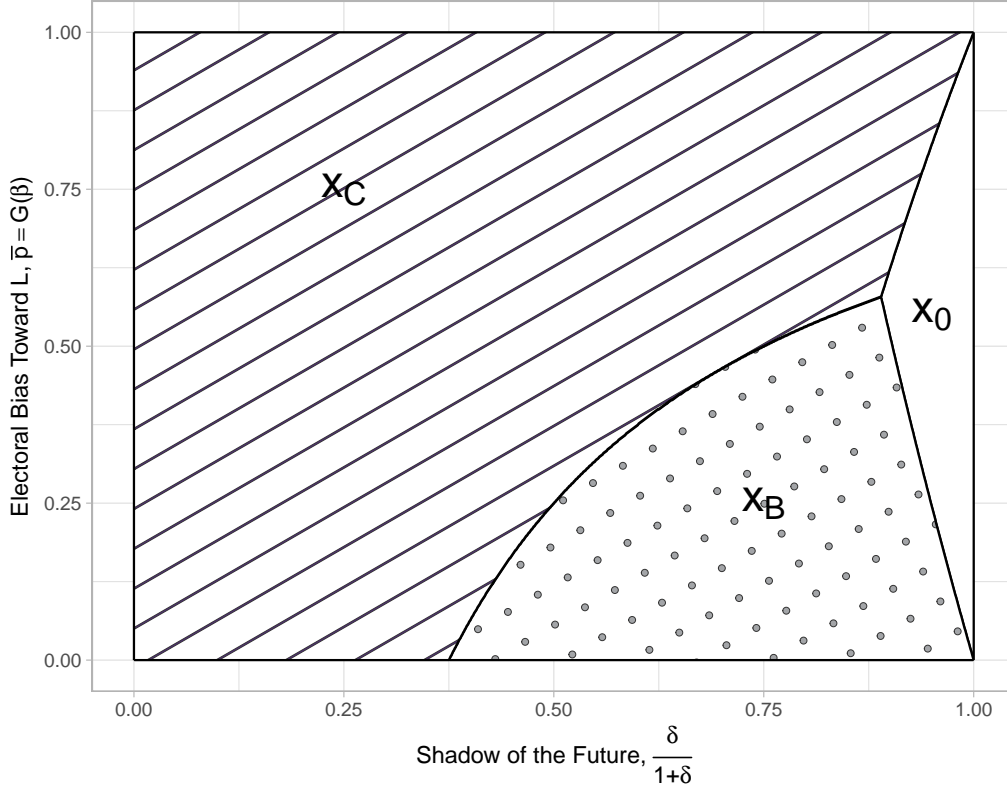


Figure 2: What Does F Offer to L ? Exogenous Election

Unfortunately for F , exploiting L 's dovishness is less desirable when the election is close (δ large), especially if R is likely to win (β small). An election victory for R leads to renegotiation to a much less desirable deal, x_A . Faced with an immediate election and significant prospects that R will win, F might well prefer either to offer x_B or no agreement, x_0 . The deal x_B is sustained whoever wins the election and is attractive especially if R is likely to win. This is shown in the grey/dotted region. Alternatively, F could form no deal and then exploit the winner as much as possible after the election (see the first column of Table 1). This option is the most preferred when the election is imminent and is shown in the white/nonshaded region.

When elections are exogenous, L is unencumbered by electoral incentives, so he will always accept what F offers. Then, for F , first period offers simply navigate a risk-return tradeoff that exploits L 's dovishness in the first period, but might yield a suboptimal deal in the second period if R wins the election. F offers the more generous agreement x_B when R is likely to win the election. If the election is

very close, then F simply prefers to wait to engage in post-electoral bargaining that maximally exploits the second period domestic leader rather than tie its hands. Here, the key takeaway is that despite the absence of electoral incentives (i.e., desire to shape policy so as to affect the election), the *prospect of renegotiation tomorrow influences how agreements are signed today*.

3.2 Dovish Electoral Incentives

We now consider cases in which a dove's electoral fortunes depend upon the agreement signed in the first period. Voters may be retrospective, assessing the incumbent based on what he delivers prior to the election, or prospective, comparing candidates based on their expected future policy returns. We consider each electoral context in turn.

Retrospective voters evaluate the incumbent by what he or she has delivered so far (Ferejohn 1986). Referring back to Figure 1, the median voter has their own preference over policy outcomes depending on their valuation of international cooperation. If voters are relatively dovish, such that they like all three proposals, then signing any deal enhances incumbent L 's reelection prospects. Formally, if $\theta_M + v_M(x_C) > 0$, then the median voter will reward the incumbent for agreeing to the deal x_C . This means that L will always accept x_C if it is offered, because it helps him electorally.

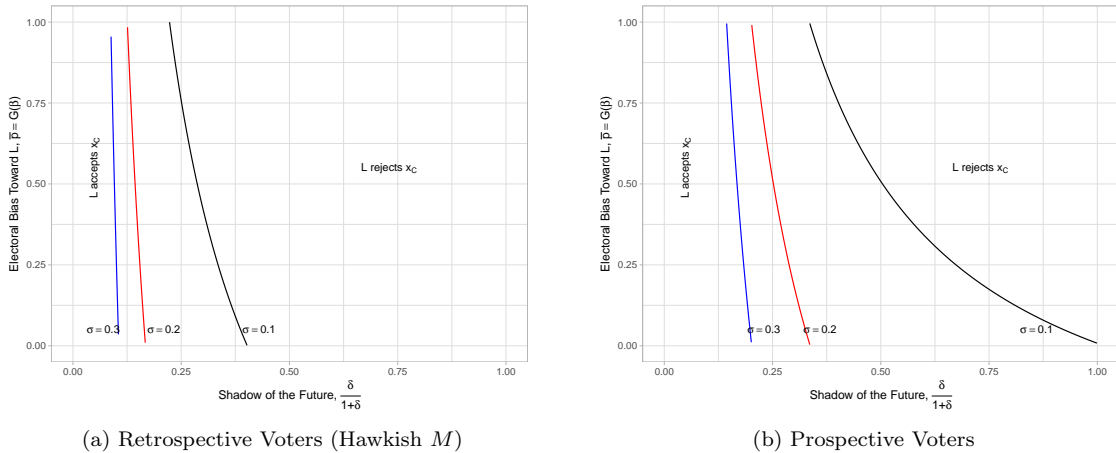


Figure 3: When Does L Accept x_C ?

However, with relatively hawkish retrospective voters or prospective voters, signing x_C harms L electorally. The likelihood that he accepts x_C depends on the sensitivity of electoral outcomes to the agreement's salience, σ , and the value of holding office, Ψ . When voters are retrospective but hawkish, then $\theta_M + v_M(x_C) < 0$, meaning their utility under the deal x_C is worse than their utility under the status quo x_0 . In this scenario, voters punish L for negotiating a deal that concedes too much policy. The left panel of Figure 3 shows regions of the parameter space where L accepts and rejects x_C based

on the value of the future, the baseline probability of winning, and the sensitivity of the election to the policy deal. As σ increases, the expression $p = \bar{p} + g\sigma(\theta_M + v_M(x_C))$ decreases, making x_C a less attractive option to L because the probability of winning reelection becomes sufficiently small relative to the probability of winning reelection without an agreement in place (which is just \bar{p}). In the extreme case where the election is almost entirely determined by the policy outcome ($\sigma \rightarrow \infty$) or the leader cares primarily about office-holding ($\Psi \rightarrow \infty$), then L always rejects x_C .

When voters are prospective, L 's proclivity toward x_C is also diminished but under a different strategic logic. L has an innate bargaining disadvantage relative to R , which translates into an electoral disadvantage when voters are forward-looking. Specifically, if L negotiates x_C in the first period, voters know that R can renegotiate to x_A after the election should she win, a much more attractive outcome for D . In this instance, L 's best option is to make negotiations in the first period that nullify this advantage, like the deal x_B . Such a deal would persist in both periods as neither leader could credibly threaten to exit the agreement. The second panel of Figure 3 shows how L 's chances of accepting x_C decrease as the election becomes more sensitive to the first round policy deal.

With hawkish retrospective voters or prospective voters, we observe that electoral pressures induce L to bargain harder, rejecting deals he would myopically accept. Generated by a desire to stay in office, L 's behavior effectively mirrors the ratification constraint articulated by Schelling (1960) and Putnam (1988). This constraint enhances a dove's bargaining leverage, not because of an institutional constraint (e.g., Milner and Rosendorff 1997), but because of the desire to keep his job. In essence, the presence of electoral incentives imposes constraints on the likelihood of L signing deals he would prefer *ex ante*.

3.3 Bargaining with Doves in the Shadow of Elections

Holding elections in D also affects how F proposes deals in the first period. Here, we investigate how the aforementioned electoral tensions that L faces shape preferences over policy prior to the election. An overarching theme of the analysis is that F wants to maximally exploit the domestic leader, or conclude an agreement in which the domestic nation signs away as much of its policy autonomy as possible. When L is the incumbent, this naturally means that F will propose x_C under a wide range of conditions.

Suppose voters are retrospective. The general setup is similar to the exogenous case illustrated above. When the election is imminent (δ large), F forgoes first period concessions and simply extracts the largest possible concession after the election, shown in the white/nonshaded regions in Figure 4. If the electoral bias favors R (β small) such that R is likely to win, then F offers x_B and such a deal persists through both rounds, shown in the grey/dotted regions.

Absent an imminent election or a large electoral bias in favor of R , F wants to maximally exploit the

dove by demanding x_C . As mentioned, the conditions under which L accepts this deal depend on the preferences of the voters. If F offers this deal, then it obtains considerable policy concessions in the first round. However, should R win the election, the deal will be renegotiated to x_A .

If the electorate is dovish, then signing the agreement x_C is beneficial to L electorally. For F , this is strategically very valuable: L is willing to accept a deal that grants F large policy concessions, and such a deal makes L more attractive electorally to voters. This also implies that L becomes more likely to be the winner of the election, and would not be able to renegotiate the terms of the deal in the second period. Hence, F 's incentives to offer the deal x_C to L when the electorate is dovish and retrospective increase over the baseline exogenous case, denoted by the expansion of the purple/striped region in the first panel of Figure 4.

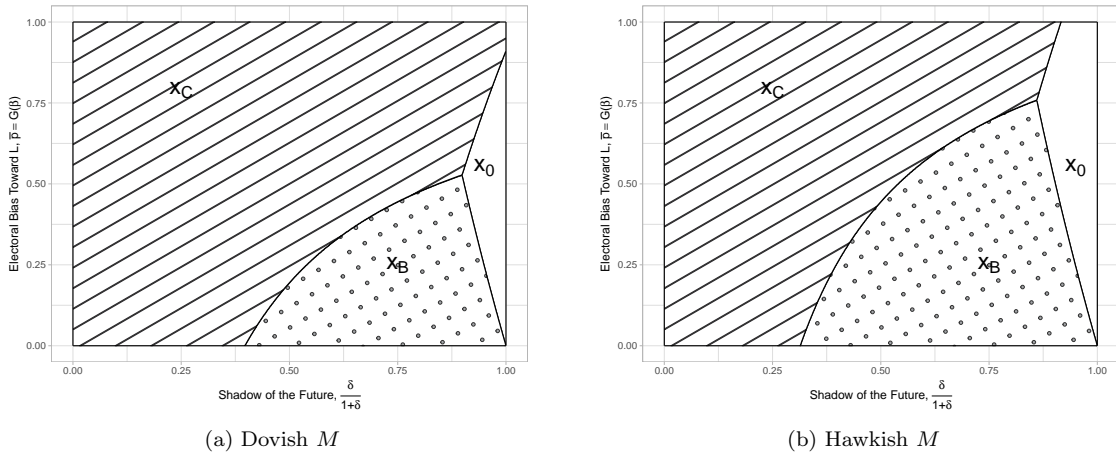


Figure 4: What Does F Offer to L ? Retrospective Election

The opposite effect obtains if the electorate is retrospective but relatively hawkish. In this case, since L would be punished for agreeing to x_C , forging such a deal in the first period increases the chances that R becomes the leader in the second period. Such an outcome also means that the second period agreement would move from x_C to x_A , F 's least preferred deal. Consequently, F is less likely to offer x_C to L , shown by the contraction of the purple/striped region in the second panel of Figure 4. We emphasize here that F 's strategic calculus is shaped by the fact that a myopically generous concession from D actually affects who F would be bargaining with in the second period. The hawkish electorate deters F from offering the exploitative deal x_C : the shadow of the election shapes F 's pre-electoral bargaining.

F 's incentives to offer x_C when voters are retrospective vary importantly with the way in which the median voter values having an agreement in place. Clearly, when θ_M increases, the purple/striped region gets larger. Since voters reward L electorally for striking this agreement, F is incentivized to demand

maximal concessions. Additionally, as the election becomes more sensitive to the agreement, increasing σ , there is a bigger difference in the expansion and contraction of this region as the median voter becomes relatively more hawkish or dovish.




Figure 5: What Does F Offer to L ? Prospective Election

Finally, we develop F 's optimal offers when voters are prospective. A prospective electorate selects a candidate by anticipating the downstream policy agreement. Voters consider who can deliver greater utility in second period negotiations and cast their votes in favor of that candidate. When L is the incumbent, we know from Table 1 that he has a bargaining disadvantage relative to R going into the second period. As such, F is deterred from offering x_C to L unless the election is far off (δ small) or there is near certainty that L will win reelection (β large). Instead, F is more likely to offer x_B , a deal that would persist in both periods and would nullify R 's bargaining advantage relative to L after the election. The growth of the grey/dotted region in Figure 5 illustrates this. In common with earlier cases, when an election is imminent, F prefers to wait and offer x_0 , the white/nonshaded region, rather than tie its hands. This strategy maintains R 's bargaining advantage and means it is likely that x_B is concluded in the second period, but should L win reelection then F demands x_C after the election.

Bargaining in the shadow of a responsive electorate shapes incentives for the domestic incumbent to accept and reject international agreements based on their likelihood of leading to electoral victory. But

the pressure of voters also extends outward: different electoral contexts in D influence how F bargains with L prior to the election.

4 Hawkish Incumbent

Different strategic dynamics arise when bargaining with a hawkish leader in the first period. F 's general goal is to maximally exploit whatever R will accept, which varies based on her electoral constraints. We believe that the most interesting case is when a hawkish leader negotiates in the shadow of prospective voters: since leaders' primary objective is to retain office, the hawk will not agree to any deal that nullifies her bargaining advantage because it reduces her electability. She can retain her bargaining advantage in one of two ways. First, she can reject deals, even those that provide nation D with large policy concessions, because the voters know that only she can deliver a harder line after the election. Second, in a more extreme case, the hawk might even agree to weak deals that grant F large concessions, knowing full well that she will renege on the deal and obtain additional concessions after securing reelection, while the dove would be stuck with the weak deal.

4.1 Pure Renegotiation Concerns with Hawks

Myopically, R is only willing to accept x_A or x_B , deals which would persist across both periods as neither domestic leader could renegotiate after the election. In a single shot game – i.e., if the policy deal were final after the first period – F would offer x_B , which R would accept. This result is consistent with other theoretical and empirical findings that hawkish leaders have greater leverage than doves to deliver better policy outcomes (Schultz 2005; Buisseret and Bernhardt 2018; Mattes and Weeks 2019).

In the exogenous election context, F can obtain the agreement x_B in both periods, and indeed, when R is the likely winner of the election this is exactly what F does. However, by offering x_B in the first period, F forgoes the opportunity to exploit L 's dovishness should L win the election. The grey/dotted region of Figure 6 shows the parameters under which F offers x_B , the deal F and R would conclude in a single-shot game.

If an election is imminent (δ large), then F forgoes a first round deal, $x_1 = x_0$, and following the election F fully exploits the winner, with the second period outcome being x_C or x_B depending on who wins the election. By waiting, F avoids tying its hands and retains the option of exploiting L 's dovishness. Yet, waiting is a poor choice when an election is far off as the surplus from cooperation goes unrealized. The range of the parameter space for which F waits is illustrated with the white/nonshaded region of Figure 6.

Although R would not myopically agree to the deal x_C , as $\theta_R + v_R(x_C) < 0$, counterintuitively

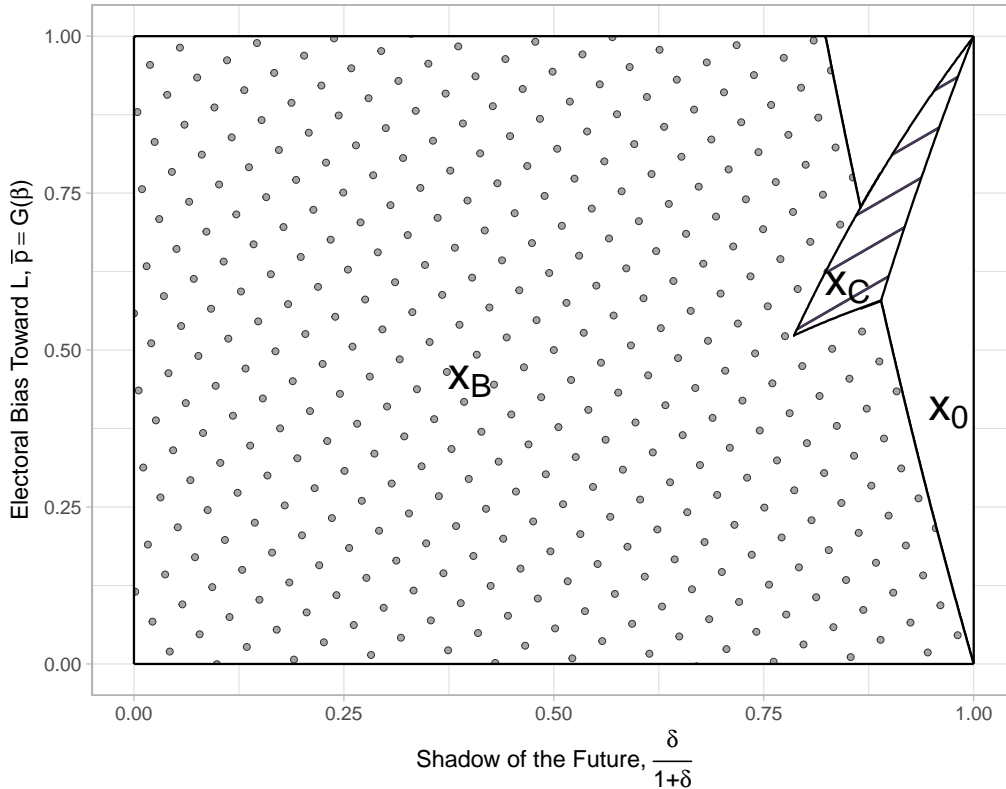


Figure 6: What Does F Offer to R ? Exogenous Election

there are conditions under which she might agree to such a deal even when the deal has no electoral consequences. The conditions under which this deal is offered and accepted is characterized by a twoway risk-return tradeoff. If R accepts this deal and L wins the election, then x_C persists and F obtains maximal policy concessions across both periods. However, should R win the election, she can obtain her most preferred agreement, x_A , in the second period. The hawk trades off a bad deal today, something she would not myopically accept, for the possible implementation of a better agreement after the election should she win. By the same token, F maximally exploits the domestic nation today, at the cost of potentially inferior second period outcomes. The intersection of these strategic calculi forms the purple/striped region in Figure 6.

When there is a hawkish incumbent, F faces an important intertemporal tradeoff. There is always the possibility that a more dovish leader comes to power after the election. If F ties its hands by offering something like x_B , then it is unable to maximally benefit from a dove's willingness to concede even more policy in future negotiations. But if F demands too much today and the hawk is reelected, then renegotiations would bring about less desirable outcomes.

Although it is possible for F to offer x_C and R to accept it, such deals require a delicate balance, so we would rarely expect to see this in practice. However, as we shall subsequently see, electoral incentives

can make this deal an attractive option for both F and R . The dominant themes in the exogenous election context is between offering the deal the hawk will agree to and having this deal persist (which is attractive when the election is far away or R is likely to win) and simply waiting until after the election and concluding a deal and maximally exploiting whoever wins (which is attractive when the election is close or L is likely to win).

4.2 Hawkish Electoral Incentives

Now suppose that R faces electoral consequences when negotiating certain deals. Again, we think about retrospective and prospective electorates.

With retrospective voters, an incumbent R only delivers policy outcomes that the median voter prefers relative to the status quo. This is always the case for x_A and x_B , $\theta_M + v_M(x_A) > 0$ and $\theta_M + v_M(x_B) > 0$. If M is relatively hawkish, meaning $\theta_M + v_M(x_C) < 0$, then R does not accept x_C . This deal is not myopically preferred by R anyway, and agreeing to it would harm her electorally. The tradeoff described by the purple/striped region of Figure 6 disappears in this electoral context. By contrast, if the median voter is relatively dovish, then R would be willing to accept x_C under a wider range of conditions relative to the exogenous baseline. Even though she myopically does not like this deal, it helps her win reelection because she delivers an agreement that the dovish voters like. What is more, R would be able to renegotiate to x_A after the election, should she win.

R 's electoral incentives are most interesting and counterintuitive when voters are prospective. In a prospective context, R has a stronger post-electoral bargaining position relative to L , which makes her more attractive to voters. This is because R can deliver a better second period outcome than L (x_A and x_B from x_C and x_0 respectively). Then, in pre-election bargaining R wants outcomes that enhance this bargaining advantage, while, as we saw, dovish L seeks outcomes that deny R a bargaining edge in the eyes of prospective voters. Strikingly, R maximizes her electoral odds by failing to conclude an agreement and allowing the status quo x_0 to persist until after the election. Given the assumption that R would have proposal power with a pre-existing agreement from which she would walk away after the election, she could also enhance her electoral odds by agreeing to x_C before the election and then obtaining x_A afterward, the best possible agreement for nation D . Conversely, R forsakes her superior electoral position if she agrees to x_B .

4.3 Bargaining with Hawks in the Shadow of Elections

R 's electoral concerns affect F 's likelihood of demanding certain policy concessions. As we just discussed, the electoral context affects which agreements R will accept and who is likely to win the election. These

factors determine what deals F offers. Notably, different electoral contexts affect R 's willingness to accept x_C , a deal she myopically dislikes.

We start by considering retrospective voters. When such voters are hawkish, R is reluctant to accept x_C , as discussed above. Consequently, F is reluctant to offer it, because it knows that such a deal would be rejected. When the election is close, F simply waits and sees who wins the election and makes a deal in the second period. But if the wait until the election is long, then F realizes the gains of cooperation by offering the deal that the hawk myopically prefers, x_B . This is shown in the right panel of Figure 7.

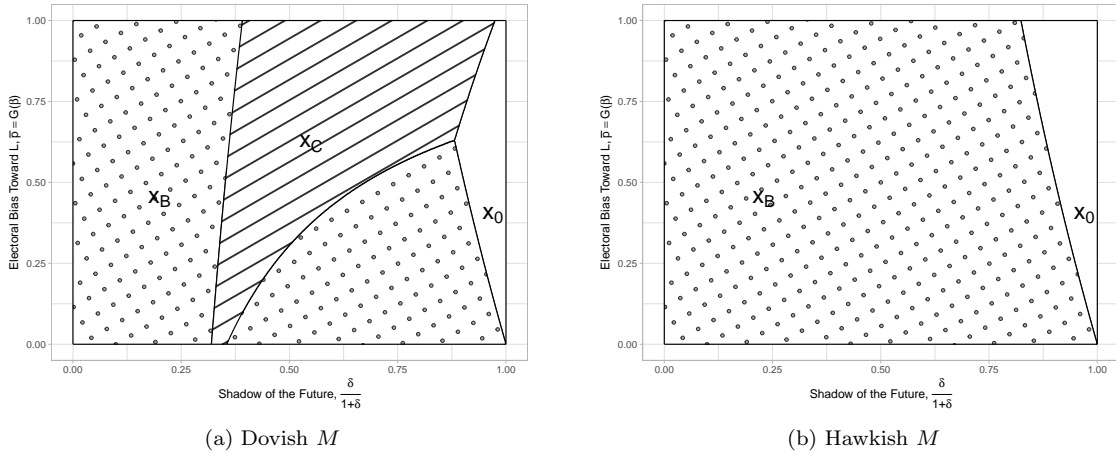


Figure 7: What Does F Offer to R ? Retrospective Election

A more interesting case occurs when the median voter is dovish and retrospective, since such voters will reward a hawk for delivering greater policy concessions. F can capitalize on its desire to maximally exploit the incumbent, even demanding x_C . Relative to the exogenous baseline, the purple/striped region in the left panel of Figure 7 is much larger, showing that R is now willing to accept x_C because it aides her reelection. Note however that F 's willingness to offer this deal obtains when L is likely to win the election (β large) or if the election is not too close (δ small enough), so as to avoid renegotiation to x_A if R were to be reelected. That is, even though R concedes more, which bolsters her reelection prospects, F may only be willing to offer such a deal if L is the likely leader in the second period. The dynamic that the hawk will accept a more dovish deal is exacerbated as officeholding becomes important ($\Psi \rightarrow \infty$), because R 's desire to retain her job outweighs her preference for a less conciliatory agreement.

We turn finally to the case when voters are prospective. If the hawk accepts the deal she would myopically prefer, x_B , then she undermines her own electoral advantage. In contrast, if she accepts x_C , the deal she myopically dislikes – which we find to be less plausible substantively – she enhances her electoral advantage because she, and not the dovish leader L , could renegotiate after the election. When electoral incentives dominate ($\Psi \rightarrow \infty$), this produces a paradoxical outcome in which the hawk only

accepts the deal she myopically dislikes, or punts negotiations until after the election. We consider this in Figure 8.

Since the hawk rejects x_B , F 's bargaining decision is reduced to whether or not it wants to strike x_C in the first period, or retain the status quo of no deal. In the purple/striped region on the left, F maximally exploits nation D in the first period, but risks the possibility of renegotiation in the future. On the right in the white/nonshaded region, as F cares more about post-electoral bargaining, it simply prefers to wait until after the election to deal with the second period leader in D .




Figure 8: What Does F Offer to R ? Prospective Election

When domestic leaders are driven by office-holding concerns, they agree to almost any terms so long as it promotes their electoral odds. This willingness to give up policy to enhance electoral prospects potentially enables F to obtain superior short-term deals from hawks, but prevents F from fully exploiting doves. If the election is reasonably distant, then F wants to maximally exploit both a hawk and dove's willingness to make deals to extract immediate concessions, although at the risk of future renegotiations. In contrast, if an election is close, often F 's best strategy is to simply wait in the hope of exploiting a dove after the election.

5 Discussion

The dynamics of real-world negotiations are complex. Our analysis centers on the interplay between elections and the political dynamics that underlie interstate bargaining. We highlight three important roles that elections have in shaping international agreements. Elections introduce uncertainty in the stability of policy: the possibility of leader turnover in the future means that a deal signed today might not persist tomorrow. Elections also affect the types of agreements that leaders are willing to sign. Finally, the shadow of elections directly influences the types of deals that are proposed in pre-electoral bargaining.

Since elections have several effects on the signing of international agreements, making broad predictions is difficult. Our analysis isolates each of the three mechanisms listed and describes which agreements become more or less attractive to domestic and foreign policymakers in the presence of different electoral contexts, incentives, and pressures. There are several key themes that run through the analysis, but some of our results are sensitive to how people assess their leaders.

The shadow of renegotiation after an election affects pre-electoral bargaining. We underscore that the leverage needed to renegotiate an agreement is only credible when one player is willing to walk away from the deal. In many bargaining models there is both an inherent first-mover advantage as well as a prediction that instantaneous,⁵ efficient outcomes are negotiated (Rubinstein 1982). These factors are not what drives renegotiation in our model. We find that the ability to renegotiate comes directly from what leaders are willing to myopically accept. For example, because a hawkish leader myopically prefers no deal to the deal x_C , she is able to renegotiate x_C if it is the status quo at the start of the second period.

One common claim in the literature is that hawks have greater bargaining leverage than doves. Our model exhibits this feature when leaders are myopic. However, when elections and renegotiations are taken into account the predictions are more nuanced. Differences in bargaining leverage of hawks and doves drive many of the results, but they often do so off the equilibrium path as leaders on both sides factor in how agreements today affect future negotiations and with whom they will have to bargain.

Introducing domestic electoral incentives changes the conditions under which L and R are willing to accept agreements that they would prefer myopically. In a retrospective electoral context, voters reward the incumbent for improving the status quo. To ingratiate themselves with the voters, doves reject deals that the voters dislike, even though F and the dove could agree to such deals on policy terms. The threat of removal by a hawkish electorate forces the dovish incumbent to bargain harder, resembling

⁵In our model, it is a perfectly rational decision to forgo cooperation today with the expectation that one can conclude more favorable terms tomorrow or that rejecting an agreement could enhance electability. While post-election outcomes are efficient on the equilibrium path, these bargaining surpluses need not be reached immediately.

the traditional ratification constraint (Schelling 1960; Putnam 1988). Office-seeking hawks are willing to make additional concessions over what they support on policy grounds in order to deliver benefits to dovish voters, for which they hope to be rewarded.

Ex ante, hawkish leaders have a bargaining advantage over dovish leaders. When the voters are prospective, this translates into an electoral advantage for hawks by either failing to conclude agreements or accepting weak agreements that the voters know only they can renegotiate in the second period. The desire to be reelected generates counterintuitive behavior on the part of a hawkish incumbent. Hawks seek to differentiate themselves from doves by rejecting any deal that undercuts their electoral prospects.

Prospective electoral incentives also cause doves to bargain harder, as failing to conclude a good agreement (x_B) hurts L electorally. Dovish leaders again are more likely to reject deals that they myopically like in order to maintain their jobs. The desires of incumbents to maximize their electoral prospects can result in leaders undermining their immediate policy successes. Paradoxically, this can even lead to doves cutting better deals than hawks.

There is no simple, reduced-form effect of elections on international negotiations. What gets negotiated depends upon who is power before the election, how close the election is, who is likely to win, and how the voters evaluate potential candidates. Our model provides a way to isolate these various mechanisms and sheds light on when deals are likely to be struck, how they affect elections, and the ability of nations to exercise their bargaining power.

6 Appendix

In this appendix, we present the model in which F can offer a discrete set of deals to the leader of D . In our supplemental materials we present a more general version of the model in a continuous policy space setting.

6.1 Second Period Agreements

Let $\hat{x}_j(x_1)$ be the second period agreement given first period agreement x_1 if leader $j \in \{L, R\}$ is elected. For L , the value of agreement x_2 in the final period is $\theta_L + v_L(x_2)$ and the value of no agreement is $v_L(x_0) = 0$. Hence L would remain in deal x_A, x_B or x_C . In contrast, R would exit agreement x_C but could not credibly exit from x_A or x_B .

Proposition 1 *If the first period agreement is x_1 and leader j is elected, then the second period outcome, $x_2 = \hat{x}_j(x_1)$, is shown by Table 1.*

Table 1: Second Period Agreement $x_2 = \hat{x}_j(x_1)$ given Leader $j \in \{L, R\}$ in Power and the First Period Agreement, x_1

First Period Agreement, x_1	x_0	x_A	x_B	x_C
L wins election	x_C	x_A	x_B	x_C
R wins election	x_B	x_A	x_B	x_A

6.2 Electoral Probabilities

Let $\bar{p} = G(\beta)$ be the baseline probability that L is elected. As a simplification, assume that density of

Table 2: Electoral Probabilities that L is Elected

Electoral Assumption	Incumbent L	Incumbent R
Exogenous elections	$\bar{p} = G(\beta)$	$\bar{p} = G(\beta)$
	Retrospective Voters	
x_0	$p_0 = G(\beta)$	$p_0 = G(\beta)$
$x_1 = x_B$	$p_B = G(\beta + \sigma(\theta_M + v_M(x_B)))$	$p_B = G(\beta - \sigma(\theta_M + v_M(x_B)))$
$x_1 = x_C$	$p_C = G(\beta + \sigma(\theta_M + v_M(x_C)))$	$p_C = G(\beta - \sigma(\theta_M + v_M(x_C)))$
	Prospective Voters	
Generic, x_1	$p_{x_1} = G(\beta + \sigma(v_M(\hat{x}_L(x_1)) - v_M(\hat{x}_R(x_1))))$	
x_0	$p_{x_0} = G(\beta + \sigma(v_M(x_C) - v_M(x_B)))$	
x_B	$p_{x_B} = G(\beta + \sigma(v_M(x_B) - v_M(x_B))) = G(\beta)$	
x_C	$p_{x_C} = G(\beta + \sigma(v_M(x_C) - v_M(x_A)))$	

G is uniform. $G(\beta + \sigma y) = G(\beta) + g\sigma y$ where $g = G'$.

6.3 Dovish Incumbent

6.3.1 Agreements that L Accepts

It is straightforward to see that L always accepts x_B : myopically he supports the agreement and such an agreement maximizes his probability of election. However, L only accepts x_C if $U_L(x_C) \geq U_L(x_0)$.

We define

$$\alpha_C^L = \begin{cases} 1 + \frac{\theta_L + v_L(x_C)}{\delta(v_L(x_A) - v_L(x_B))} + \frac{g\sigma(v_M(x_C) + \theta_M)(-v_L(x_A) + v_L(x_C) + \Psi)}{v_L(x_A) - v_L(x_B)} & \text{if retrospective election} \\ 1 + \frac{\theta_L + v_L(x_C)}{\delta(v_L(x_A) - v_L(x_B))} - \frac{g\sigma\Psi(v_M(x_A) - v_M(x_B))}{v_L(x_A) - v_L(x_B)} \\ + \frac{g\sigma((v_L(x_A) - v_L(x_C))v_M(x_A) - (v_L(x_B) - v_L(x_C))v_M(x_B) - (v_L(x_A) - v_L(x_B))v_M(x_C))}{v_L(x_A) - v_L(x_B)} & \text{if prospective election} \end{cases} \quad (1)$$

The exogenous election case corresponds to $\sigma = 0$ (removes the final term in each case). If $\bar{p} = G(\beta) \leq \alpha_C^L$ then $U_L(x_C) \geq U_L(x_0)$ and so L accepts a first period offer of x_C . In the exogenous election, or the retrospective election with a dovish median voter ($\theta_M + v_M(x_C) > 0$), then $\alpha_C^L > 1$ so L always accepts x_C . If the election context is retrospective and the median voter is hawkish ($\theta_M + v_M(x_C) < v_M(x_0)$) or if the election context is prospective, then as office holding incentives dominate ($\Psi \rightarrow \infty$) L always rejects x_C .

6.3.2 F 's Preferred Offer to L

Define ρ_{C0}^L as the probability $\bar{p} = G(\beta)$ such that F is indifferent between the first round deals x_C and x_0 . If $\bar{p} = G(\beta) \geq \rho_{C0}^L$, then F prefers x_C to x_0 . Define ρ_{B0}^L and ρ_{BC}^L as the analogous indifferences from F 's perspective between x_B and x_0 and between x_B and x_C . If $\bar{p} = G(\beta) \geq \rho_{BC}^L$, then F prefers the first round outcome x_C to the first round outcome x_B . If $\bar{p} = G(\beta) \geq \rho_{B0}^L$, then F prefers the first round outcome x_0 to the first round outcome x_B .

$$\rho_{C0}^L = \begin{cases} 1 - \frac{\theta_F + v_F(x_C) - v_F(x_0)}{\delta(v_F(x_B) - v_F(x_A))} - \frac{g\sigma(\theta_M + v_M(x_C))(v_F(x_C) - v_F(x_A))}{v_F(x_B) - v_F(x_A)} & \text{if retrospective} \\ 1 + \frac{\theta_F - v_F(x_0) + v_F(x_C)}{\delta v_F(x_A) - \delta v_F(x_B)} \\ + \frac{g\sigma(v_F(x_C)(-v_M(x_A) + v_M(x_B)) + v_F(x_A)(v_M(x_A) - v_M(x_C)) + v_F(x_B)(-v_M(x_B) + v_M(x_C)))}{v_F(x_A) - v_F(x_B)} & \text{if prospective} \end{cases} \quad (2)$$

$$\rho_{B0}^L = \begin{cases} \frac{\theta_F + v_F(x_B) - v_F(x_0)}{\delta v_F(x_C) - \delta v_F(x_B)} & \text{if retrospective} \\ -\frac{\theta_F - v_F(x_0) + v_F(x_B)}{\delta v_F(x_B) - \delta v_F(x_C)} + g\sigma(v_M(x_B) - v_M(x_C)) & \text{if prospective} \end{cases} \quad (3)$$

$$\rho_{BC}^L = \begin{cases} -\frac{v_F(x_C) + \delta v_F(x_A) - (1 + \delta)v_F(x_B)}{\delta(v_F(x_C) - v_F(x_A))} - g\sigma(\theta_M + v_M(x_C)) & \text{if retrospective} \\ \frac{\delta v_F(x_A) - (1 + \delta)v_F(x_B) + v_F(x_C)}{\delta(v_F(x_A) - v_F(x_C))} + g\sigma(v_M(x_A) - v_M(x_C)) & \text{if prospective} \end{cases} \quad (4)$$

The exogenous election conditions are given when $\sigma = 0$. Given the definitions α_C^L , ρ_{C0}^L , ρ_{B0}^L and ρ_{BC}^L , we can characterize subgame perfect equilibrium outcomes. For instance, x_C is the SPE outcome when L will accept x_C , that is to say $\bar{p} \leq \alpha_C^L$, and F prefers the outcome x_C to x_B and x_0 ($\bar{p} \geq \rho_{BC}^L$ and $\bar{p} \geq \rho_{C0}^L$). The proposition follows directly from simple logical statements about which of the acceptable proposals F most prefers:

Proposition 2 *If L is the incumbent, then the first period deal is*

$$x_1 = \begin{cases} x_B & \text{if } \bar{p} \leq \rho_{B0}^L \text{ and (either } \bar{p} \leq \rho_{BC}^L \text{ or } \bar{p} \geq \alpha_C^L) \\ x_C & \text{if } \bar{p} \leq \alpha_{LC0} \text{ and } \bar{p} \geq \rho_{C0}^L \text{ and } \bar{p} \geq \rho_{BC}^L \\ x_0 & \text{if } (\bar{p} \geq \alpha_{LC0} \text{ or } \bar{p} \leq \rho_{C0}^L) \text{ and } \bar{p} \geq \rho_{B0}^L \end{cases} \quad (5)$$

These conditions define the thresholds between the regions shown in the figures throughout the paper. We next examine how the thresholds shift in response to changes in different parameters.

Proposition 3 *The sign of the comparative statics of how the thresholds change with respect to the parameters are given in Table 3 for both retrospective and prospective election settings. For threshold y and parameter z , the table provides the sign of $\frac{dy}{dz}$. The hawk and dove references refers to the median voter (specifically, $\theta_M + v_M(x_C) < 0$ implies hawk in the retrospective context).*

Table 3: Comparative Statics for the thresholds with dovish incumbent (L) with respect to Ψ , σ , θ_M and δ .

Threshold/parameter	Retrospective Election				Prospective Election			
	Ψ	σ	θ_M	δ	Ψ	σ	θ_M	δ
α_C^L	-	-	+	-	-	-	0	-
ρ_{C0}^L	0	+hawk, -dove	-	+	0	?, depends on parameters	0	+
ρ_{B0}^L	0	0	0	-	0	+	0	-
ρ_{BC}^L	0	+hawk, -dove	-	+	0	+	0	+

For a prospective election and a retrospective election with hawkish voters: As Ψ , σ , or δ increase, L is more likely to reject x_C . In the retrospective case (with hawkish voters), increases in θ_M increases the likelihood that L accepts x_C .

The value of office holding for D does not affect F 's preferences over deals, although increases in Ψ reduce the parameters for which L will accept x_C .

6.4 Hawkish Incumbent

6.4.1 Agreements that R Accepts

Analogous to the approach above, define α_C^R such that if $\bar{p} \leq \alpha_C^R$ then R prefers the first period agreement x_C rather than waiting (x_0) i.e. \bar{p} such that $U_R(x_C) \geq U_R(x_0)$. Likewise define α_B^R such that if $\bar{p} \geq \alpha_B^R$ then R prefers the first period agreement x_B rather than waiting (x_0) i.e. \bar{p} such that $U_R(x_B) \geq U_R(x_0)$. This condition is always satisfied in the retrospective election context.

$$\alpha_C^R = \begin{cases} 1 + \frac{\theta_R + v_R(x_C) - v_R(x_0)}{\delta(v_R(x_A) - v_R(x_B))} + \frac{g\sigma(\theta_M + v_M(x_C))(\Psi + v_R(x_A) - v_R(x_C))}{v_R(x_A) - v_R(x_B)} & \text{if retrospective election} \\ 1 + \frac{\theta_R - v_R(x_0) + v_R(x_C)}{\delta(v_R(x_A) - v_R(x_B))} & \\ -g\sigma v_M(x_C) + \frac{g\sigma(v_M(x_A)(\Psi + v_R(x_A) - v_R(x_C)) - v_M(x_B)(\Psi + v_R(x_B) - v_R(x_C)))}{v_R(x_A) - v_R(x_B)} & \text{if prospective election} \end{cases} \quad (6)$$

$$\alpha_B^R = \begin{cases} 0 & \text{if retrospective election} \\ \frac{-\theta_R + v_R(x_0) - v_R(x_B)}{\delta(v_R(x_B) - v_R(x_C))} + \frac{g\sigma(v_M(x_B) - v_M(x_C))(\Psi + v_R(x_B) - v_R(x_C))}{v_R(x_B) - v_R(x_C)} & \text{if prospective election} \end{cases} \quad (7)$$

6.4.2 F 's Preferred Offer to R

We define ρ_{B0}^R as the value of $\bar{p} = G(\beta)$ such that F is indifferent between the first period outcome x_B and no first period agreement, x_0 . Likewise ρ_{C0}^R and ρ_{BC}^R define F 's indifference between x_C and x_0 and x_B and x_C , respectively.

$$\rho_{B0}^R = \begin{cases} \frac{\theta_F - v_F(x_0) + v_F(x_B)}{\delta(v_F(x_C) - v_F(x_B))} & \text{if retrospective election} \\ \frac{\theta_F - v_F(x_0) + v_F(x_B)}{\delta(v_F(x_C) - v_F(x_B))} + g\sigma(v_M(x_B) - v_M(x_C)) & \text{if prospective election} \end{cases} \quad (8)$$

$$\rho_{C0}^R = \begin{cases} 1 + \frac{\theta_F - v_F(x_0) + v_F(x_C)}{\delta(v_F(x_A) - v_F(x_B))} - \frac{g\sigma(v_F(x_A) - v_F(x_C))(\theta_M + v_M(x_C))}{v_F(x_B) - v_F(x_A)} & \text{if retrospective election} \\ 1 - \frac{\theta_F - v_F(x_0) + v_F(x_C)}{\delta(v_F(x_B) - v_F(x_A))} & \\ + \frac{g\sigma(v_F(x_C)(v_M(x_A) - v_M(x_B)) - v_F(x_A)(v_M(x_A) - v_M(x_C)) + v_F(x_B)(v_M(x_B) - v_M(x_C)))}{v_F(x_B) - v_F(x_A)} & \text{if prospective election} \end{cases} \quad (9)$$

$$\rho_{BC}^R = \begin{cases} \frac{(1+\delta)v_F(x_B) - \delta v_F(x_A) - v_F(x_C)}{\delta(v_F(x_C) - v_F(x_A))} + g\sigma(\theta_M + v_M(x_C)) & \text{if retrospective election} \\ \frac{-\delta v_F(x_A) + (\delta+1)v_F(x_B) - v_F(x_C)}{\delta(v_F(x_C) - v_F(x_A))} + g\sigma(v_M(x_A) - v_M(x_C)) & \text{if prospective election} \end{cases} \quad (10)$$

Given the definitions α_C^R , α_B^R , ρ_{B0}^R , ρ_{C0}^R and ρ_{BC}^R , we can characterize subgame perfect equilibrium outcomes. For instance, x_C is the SPE outcome when R will accept x_C , that is to say $\bar{p} \leq \alpha_C^R$ and F prefers the outcome x_C to x_0 ($\bar{p} \geq \rho_{C0}^R$) and (either F prefers x_C to x_B ($\bar{p} > \rho_{BC}^R$) or R rejects x_B ($\bar{p} < \alpha_B^R$)). The proposition follows directly from simple logical statements about which of the acceptable

proposals F most prefers:

Proposition 4 *If R is the incumbent, then the first period deal is*

$$x_1 = \begin{cases} x_C & \text{if } \bar{p} \leq \alpha_C^R \text{ and } \bar{p} \geq \rho_{C0}^R \text{ and (either } \bar{p} \geq \rho_{BC}^R \text{ or } \bar{p} \leq \alpha_B^R) \\ x_B & \text{if } \bar{p} \geq \alpha_B^R \text{ (which is always true in the retrospective setting or as } \Psi \rightarrow 0) \\ & \text{and } \bar{p} \leq \rho_{B0}^R \text{ and (either } \bar{p} \leq \rho_{BC}^R \text{ or } \bar{p} \geq \alpha_C^R) \\ x_0 & \text{if (either } \bar{p} \geq \rho_{B0}^R \text{ or } \bar{p} \leq \alpha_B^R) \text{ and (either } \bar{p} \leq \rho_{C0}^R \text{ or } \bar{p} \geq \alpha_C^R) \end{cases} \quad (11)$$

These conditions define the thresholds between the regions shown in the figures throughout the paper.

We next examine how the thresholds shift in response to parameters.

Proposition 5 *The sign of the comparative statics of how the thresholds change with respect to the parameters are given in Table 4 for both retrospective and prospective election settings. For threshold y and parameter z , the table provides the sign of $\frac{dy}{dz}$. The hawk and dove references refers to the median voter (specifically, $\theta_M + v_M(x_C) < 0$ implies hawk).*

Table 4: Comparative Statics for the thresholds with hawkish incumbent (R) with respect to Ψ , σ , θ_M and δ .

Threshold/parameter	Retrospective Election				Prospective Election			
	Ψ	σ	θ_M	δ	Ψ	σ	θ_M	δ
ρ_{B0}^R	0	0	0	-	0	+	0	-
ρ_{C0}^R	0	-hawk, +dove	+	+	0	?, depends on parameters	0	+
ρ_{B0}^R	0	-hawk, +dove	+	+	0	+	0	+
α_C^R	-hawk, +dove	-hawk, +dove	+	+	+	+	0	+
α_B^R					+	+	0	+

References

Battaglini, Marco and Bård Harstad. 2020. "The Political Economy of Weak Treaties." *Journal of Political Economy* 128(2):544–590. Publisher: University of Chicago.

URL: <http://proxy.library.nyu.edu/login?url=https://search.ebscohost.com/login.aspx?direct=true&db=bth&AN=1415461>
live

Brown, Joseph M. and Johannes Urpelainen. 2015. "Picking Treaties, Picking Winners: International Treaty Negotiations and the Strategic Mobilization of Domestic Interests." *Journal of Conflict Resolution* 59(6):1043–1073. Publisher: SAGE Publications Inc.

URL: <https://doi.org/10.1177/0022002713520479>

Buisseret, Peter and Dan Bernhardt. 2018. "Reelection and Renegotiation: International Agreements in the Shadow of the Polls." *American Political Science Review* 112(4):1016–1035. Publisher: Cambridge University Press.

URL: <http://www.cambridge.org/core/journals/american-political-science-review/article/reelection-and-renegotiation-international-agreements-in-the-shadow-of-the-polls/25D7590066D45B5767D9EB24DFD6A1F4>

Böhmelt, Tobias. 2019. "International Treaty Ratification and Leader Turnover." *Foreign Policy Analysis* 15(2):187–204. Publisher: Oxford Academic.

URL: <http://academic.oup.com/fpa/article/15/2/187/4959060>

Chapman, Terrence L, Johannes Urpelainen and Scott Wolford. 2013. "International bargaining, endogenous domestic constraints, and democratic accountability." *Journal of Theoretical Politics* 25(2):260–283. Publisher: SAGE Publications Ltd.

URL: <https://doi.org/10.1177/0951629812458042>

Clare, Joe. 2014. "Hawks, Doves, and International Cooperation." *Journal of Conflict Resolution* 58(7):1311–1337.

URL: <http://journals.sagepub.com/doi/10.1177/0022002713498705>

Colaresi, Michael. 2004. "When Doves Cry: International Rivalry, Unreciprocated Cooperation, and Leadership Turnover." *American Journal of Political Science* 48(3):555–570. Publisher: [Midwest Political Science Association, Wiley].

URL: <http://www.jstor.org/stable/1519916>

Cowen, Tyler and Daniel Sutter. 1998. "Why Only Nixon Could Go to China." *Public Choice* 97(4):605–

15. Publisher: Springer.
URL: <https://econpapers.repec.org/article/kappubcho/v3a973ay3a19983ai3a43ap3a605-15.htm>
- Cukierman, Alex and Mariano Tommasi. 1998. "When Does It Take a Nixon to Go to China?" *The American Economic Review* 88(1):180–197. Publisher: American Economic Association.
URL: <http://www.jstor.org/stable/116824>
- Ferejohn, John. 1986. "Incumbent Performance and Electoral Control." *Public Choice* 50(1/3):5–25. Publisher: Springer.
URL: <http://www.jstor.org/stable/30024650>
- Gadarian, Shana Kushner. 2010. "Foreign Policy at the Ballot Box: How Citizens Use Foreign Policy to Judge and Choose Candidates." *The Journal of Politics* 72(4):1046–1062. Publisher: [The University of Chicago Press, Southern Political Science Association].
URL: <http://www.jstor.org/stable/10.1017/s0022381610000526>
- Guisinger, Alexandra. 2009. "Determining Trade Policy: Do Voters Hold Politicians Accountable?" *International Organization* 63(3):533–557. Publisher: Cambridge University Press.
URL: <https://www.cambridge.org/core/journals/international-organization/article/determining-trade-policy-do-voters-hold-politicians-accountable/C3B9D8418F3B2681136E25748301BEAF>
- Healy, Andrew and Neil Malhotra. 2013. "Retrospective Voting Reconsidered." *Annual Review of Political Science* 16(1):285–306. Publisher: Annual Reviews.
URL: <http://www.annualreviews.org/doi/10.1146/annurev-polisci-032211-212920>
- Horowitz, Michael C., Allan C. Stam and Cali M. Ellis. 2015. *Why Leaders Fight*. New York: Cambridge University Press.
- Iida, Keisuke. 1993. "When and How Do Domestic Constraints Matter? Two-Level Games with Uncertainty." *The Journal of Conflict Resolution* 37(3):403–426. Publisher: Sage Publications, Inc.
URL: <http://www.jstor.org/stable/174262>
- Iida, Keisuke. 1996. "Involuntary defection in two-level games." *Public Choice* 89(3):283–303.
URL: <https://doi.org/10.1007/BF00159360>
- Lanoue, David J. 1994. "Retrospective and Prospective Voting in Presidential-Year Elections." *Political Research Quarterly* 47(1):193–205. Publisher: SAGE Publications Inc.
URL: <https://doi.org/10.1177/106591299404700110>

- Mattes, Michaela and Jessica L. P. Weeks. 2019. "Hawks, Doves, and Peace: An Experimental Approach." *American Journal of Political Science* 63(1):53–66. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12392>.
URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/ajps.12392>
- Milner, Helen V. 1997. *Interests, Institutions, and Information*. Princeton: Princeton University Press.
- Milner, Helen V. and B. Peter Rosendorff. 1997. "Democratic Politics and International Trade Negotiations: Elections and Divided Government As Constraints on Trade Liberalization." *Journal of Conflict Resolution* 41(1):117–146. Publisher: SAGE Publications Inc.
URL: <https://doi.org/10.1177/0022002797041001006>
- Mo, Jongryn. 1994. "The Logic of Two-Level Games with Endogenous Domestic Coalitions." *The Journal of Conflict Resolution* 38(3):402–422. Publisher: Sage Publications, Inc.
URL: <http://www.jstor.org/stable/174362>
- Palmer, Glenn, Tamar London and Patrick Regan. 2004. "What's Stopping You?: The Sources of Political Constraints on International Conflict Behavior in Parliamentary Democracies." *International Interactions* 30(1):1–24. Publisher: Routledge eprint: <https://doi.org/10.1080/725289044>.
URL: <https://doi.org/10.1080/725289044>
- Putnam, Robert D. 1988. "Diplomacy and Domestic Politics: The Logic of Two-Level Games." *International Organization* 42(3):427–460.
- Ramsay, Kristopher W. 2004. "Politics at the Water's Edge: Crisis Bargaining and Electoral Competition." *Journal of Conflict Resolution* 48(4):459–486. Publisher: SAGE Publications Inc.
URL: <https://doi.org/10.1177/0022002704266156>
- Reed, W. Robert and Joonmo Cho. 1998. "A Comparison of Prospective and Retrospective Voting with Heterogeneous Politicians." *Public Choice* 96(1/2):93–116. Publisher: Springer.
URL: <http://www.jstor.org/stable/30024380>
- Rubinstein, Ariel. 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50(1):97–109. Publisher: [Wiley, Econometric Society].
URL: <http://www.jstor.org/stable/1912531>
- Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge: Harvard University Press.
- Schultz, Kenneth A. 1998. "Domestic Opposition and Signaling in International Crises." *The American Political Science Review* 92(4):829–844. Publisher: [American Political Science Association, Cambridge

University Press].

URL: <http://www.jstor.org/stable/2586306>

Schultz, Kenneth A. 2005. "The Politics of Risking Peace: Do Hawks or Doves Deliver the Olive Branch?" *International Organization* 59(1):1–38. Publisher: Cambridge University Press.

URL: <http://www.cambridge.org/core/journals/international-organization/article/politics-of-risking-peace-do-hawks-or-doves-deliver-the-olive-branch/139CE70434CDBD4D8F4D6CB9006AA045>

Smith, Alastair. 2009. "Political Groups, Leader Change, and the Pattern of International Cooperation." *Journal of Conflict Resolution* 53(6):853–877. Publisher: SAGE Publications Inc.

URL: <https://doi.org/10.1177/0022002709344419>

Smith, Alastair and David R. Hayes. 1997. "The shadow of the polls: Electoral effects on international agreements." *International Interactions* 23(1):79–108. Publisher: Routledge eprint: <https://doi.org/10.1080/03050629708434902>.

URL: <https://doi.org/10.1080/03050629708434902>

Tarar, Ahmer. 2001. "International Bargaining with Two-Sided Domestic Constraints - Ahmer Tarar, 2001." *Journal of Conflict Resolution* 45(3):320–340.

URL: <http://journals.sagepub.com/doi/10.1177/0022002701045003004>

Wolford, Scott. 2012. "Incumbents, successors, and crisis bargaining: Leadership turnover as a commitment problem." *Journal of Peace Research* 49(4):517–530. Publisher: SAGE PublicationsSage UK: London, England.

URL: <https://journals.sagepub.com/doi/10.1177/0022343312440807>