Quid Pro Quo Diplomacy

Matt Malis[*], Alastair Smith[†]

New York University

June 4, 2021

Abstract

Political leaders value public demonstrations of support from foreign leaders, and frequently make concessions in order to obtain them. We model the bargaining dynamics surrounding these exchanges and their impact on the recipient leader's political survival, with a focus on top-level diplomatic visits as a means of signaling international support. Our model addresses two interrelated questions: first, we consider how symbolic displays of support from one leader to another can be informative even when they are "purchased" with concessions; and second, we derive the equilibrium price and political impact of a visit under different bargaining protocols. The incentive to make a concession in exchange for a visit generally undermines a visit's signaling value. We identify a diplomatic resource curse, where the existence of opportunities for diplomatic exchange can force leaders into accepting visit-for-concession deals that leave them worse off than if they were diplomatically isolated. Visits never occur when negotiations are fully transparent. Mutually beneficial quid pro quo diplomacy requires opacity in negotiations.

[*]Email: Malis@nyu.edu. Phone: 443-377-0762. Fax: 212-995-4184

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

Heard from White House – assuming President Z convinces trump he will investigate / "get to the bottom of what happened" in 2016, we will nail down date for visit to Washington. Good luck!

—US Special Envoy to Ukraine Kurt Volker, in text message to Ukrainian Presidential Aide Andriy Yirmak, July 25, 2019[1]

The story is by now familiar: President Trump, in pursuit of dirt on a political rival, withheld U.S. support to foreign-occupied Ukraine in order to induce President Volodymyr Zelensky to announce an investigation into the family of then-candidate and former Vice President Joe Biden. Among the forms of support withheld was a White House meeting between the two leaders. For Zelensky, the visit was no small matter: as Deputy Assistant Secretary of State George Kent relayed in the impeachment hearings arising from this exchange, leaders like Zelensky "see a meeting with the U.S. President in the Oval Office at the White House as the ultimate sign of endorsement and support from the United States" (Permanent Select Committee on Intelligence 2019). Eager for such a public demonstration of diplomatic backing, Zelensky acceded to Trump's demand, until a whistleblower from within the U.S. government intervened.

Zelensky's willingness to proffer a concession for a White House visit seems reasonable at first glance, but upon further inspection a puzzle arises. The value of the visit, according to Kent, was essentially symbolic. But if the Trump administration was selling the White House visit at a price, what exactly would be signaled by the occurrence of the visit—other than Zelensky's willingness to pay for it? The credibility of a signal generally depends on its costliness; if the cost of a signal is subsidized, its credibility is undermined. How can a signal of international support be informative when the signal is effectively purchased by its beneficiary?

Recent research has begun to examine the strategic incentives underlying public diplomacy, with a primary focus on the signaling value of top-level diplomatic exchanges and the information that they convey to an international audience (McManus 2018; Matush

[1] Savage and Williams (2019)

2020; Malis and Smith 2021). Absent from these accounts, however, is consideration of the bargaining that surrounds the exchange, and how those bargaining dynamics affect the visit's signaling value. In public diplomacy, as in all facets of political life, favors are seldom given out freely. To understand the causes and consequences of diplomatic signals, it is essential to understand the bargaining and deal-making that accompany them.

This paper presents a formal model of bargaining over diplomatic signals of support. We focus on top-level diplomatic visits as a prominent means of signaling support, and we examine the consequences of these exchanges on the recipient leader's political survival. The analysis addresses two central questions: first, how can symbolic displays of support from one leader to another be informative even when they are "purchased" with concessions; and second, what is the equilibrium price, and consequently the political impact, of a visit under varying bargaining protocols.

Our model features three players: a foreign leader, a domestic incumbent, and a domestic challenger. The foreign leader and domestic incumbent negotiate over the size of a concession that the domestic leader will provide in exchange for a diplomatic visit. Visiting carries some cost for the foreign leader, and the benefit he enjoys from the visit is conditional on the domestic leader being sufficiently secure in office as to be able to deliver the agreed-upon concession. The foreign leader's decision to visit thus signals his private assessment of the incumbent's strength. Conversely, the domestic challenger is incentivized to challenge only sufficiently weak incumbents. Upon observing the occurrence (or absence) of a diplomatic visit, the challenger updates her belief of the probability that a challenge will succeed, and is deterred from participating in one (or encouraged to do so). Importantly, the informativeness of the visit depends, among other things, on the size of the negotiated concession: the prospect of a larger (conditional) benefit makes the foreign leader more willing to visit weaker incumbents, who carry a greater risk of "defaulting" on the deal—whereas a visit granted in exchange for a smaller concession reveals an especially high degree of confidence that the incumbent will be capable of delivering. Thus in bargaining over the concession, the domestic leader must balance the increased probability of obtaining a visit, against the diminished strength of the visit's signal, as well as the direct costs of the concession.

We vary the bargaining protocol along two dimensions: first, whether the domestic or foreign leader has proposal power; and second, whether the bargaining is "open" or "closed" to the domestic challenger. Consistent with standard bargaining results, we find that each leader obtains a better a deal when she or he has proposal power. Of greater interest are the implications of bargaining transparency. Our first main result is to show that, under fully open bargaining, the domestic leader never stands to benefit from offering a concession that induces a positive probability of a visit occurring (even when she has proposal power). The incumbent is better off leaving the challenger with her prior belief, rather than facing the lottery over good and bad signals that would result from a non-degenerate bargaining outcome.

The open-door bargaining protocol provides a useful analytical benchmark, but relies on the substantively questionable assumption that a third party to a diplomatic negotiation can fully observe the bargaining process. In the latter half of the analysis we relax this assumption, and instead assume that bargaining occurs behind closed doors. In this opaque setting the challenger does not observe negotiations, but does observe whether or not a visit occurs; and, should a visit occur, he probabilistically observes the size of the concession granted. This opacity creates the possibility for quid pro quo diplomacy to occur in equilibrium. These exchanges always benefit the foreign power; their utility to the domestic incumbent, however, depends on the domestic challenger's prior belief of the likelihood of salient opportunities for concessions. When prior expectations of concession salience are high, the incumbent faces a "diplomatic resource curse": she is forced into making a large concession in order to induce a visit, because she cannot commit to having *not* made such a concession—and this leaves her worse off than if she were fully diplomatically isolated. In contrast, when it seems ex-ante unlikely that the foreign leader is interested in a concession from the incumbent—and thus unlikely that a visit will occur—quid pro quo diplomacy can be mutually beneficial. In this case, the challenger draws a less-negative inference from the absence of a visit, so the incumbent can afford to turn down unfavorable deals, and only offers a concession when doing so will improve her survival prospects. We conclude the analysis by examining comparative statics of the equilibrium price and political impact of the visit with respect to various attributes of the two leaders in the exchange and features of the bargaining protocol.

This paper relates most directly to formal models of diplomatic visits by Malis and

Smith (2019, 2021) and by Matush (2020)—which are, to our knowledge, the only game-theoretic approaches to the topic in extant literature—as well as related empirical research (Goldsmith, Horiuchi and Matush 2020; Ostrander and Rider 2019; McManus 2018; Lebovic and Saunders 2016; Nitsch 2007). More broadly, our study contributes to a number of other formal literatures as well: our theory relates to models of bargaining in front of audiences (Groseclose and McCarty 2001; Stasavage 2004; Perlroth 2019); to models of international bargaining over policy concessions (Andersen, Harr and Tarp 2006; Vreeland and Dreher 2014); and to models of informational channels through which international actors can influence domestic politics (Fang 2008; Shadmehr and Boleslavsky Forthcoming). Drawing on insights from these diverse literatures, our model examines an international bargaining process in which a material concession is exchanged for the revelation of information to a third party. This setup gives rise to some novel strategic considerations and helps makes sense of a broad swath of heretofore unexplained political activity.

1 Bargaining Over Symbolic Support

The bargaining over a prospective Trump-Zelensky visit represents, in unusually stark terms, a recurrent pattern in American diplomatic practice of exchanging concessions for diplomatic signals. Here we present a few more historical cases, and consider the insights from existing theoretical and empirical literature which can be leveraged to develop our model of bargaining over symbolic support.

A prospective visit between President Obama and Azerbaijani President Ilham Aliyev unfolded in a manner not entirely dissimilar to the Trump-Zelensky exchange depicted above. Aliyev was slotted to attend a multilateral nuclear security summit in Washington in March 2016, but hoped to leverage the opportunity to elicit an even stronger signal for his domestic audience, in the form of a one-on-one meeting with Obama. According to an Azerbaijani journalist and human rights activist, Aliyev was "eager for that ultimate seal of approval – a few minutes and a photo op with Obama – that would give him the image boost he seeks in the midst of an economic crisis at home" (Huseynov 2016). The Obama administration, aware of the value Aliyev placed on such a signal, demanded a concession in return: the release of political prisoners who were part of the focus of a broader human rights campaign. One US official involved in the negotiations noted that the deal was "made pretty darn explicitly. It was something like, 'We need the following things to happen ... There's a chance you might get to meet with the President"'.[2] Aliyev released two political prisoners, which turned out to be insufficient to earn him a meeting with the President. He instead received a one-on-one with the Vice President (Office of the Vice President 2016); a month later he effectively revoked the prior concession by imprisoning two other activists on trumped-up charges (Gogia 2020).

Such a transactional approach to the granting or withholding of diplomatic visits can be observed throughout recent American history. Seeking UNSC authorization for a military intervention in Libya earlier in his tenure, Obama turned to Gabonese President Ali Bongo Ondimba for a critical supportive vote. Bongo delivered, and in return was granted a stay at Obama's private guest residence later that spring (O'Grady 2016). When President Bush visited Poland shortly after the invasion of Iraq in 2003, "the point of his visit [was] obvious: to thank this country for supporting American policy" and to "signal ... that Poland, in the enthusiastic eyes of Washington, has become an important ally, even a special friend" (Bernstein 2003). Despite the fact that a quid pro quo was "obvious", the occurrence of the visit nonetheless carried some signaling value. Discussing the possibility of a US-Korean visit in 1964, a telegram from the US embassy in Korea advised that "[t]iming of visit should be related to progress [on] ROK-Japan normalization", an issue the US had been pushing despite domestic difficulties faced by the Korean government.[3] A 1955 telegram from Embassy Cairo likewise recommended delaying a visit with Nasser until "pendulum in Egyptian-United States relations could by other means be started again toward United States."[4] Generalizing beyond these individual anecdotes, Malis and Smith (2021) provide large-N evidence that post-war US presidents have systematically reaped concessions from the leaders with whom they conduct diplomatic visits, in the form of closer voting alignment in the UN General Assembly and increased market access for US exporters.

While the discussion thus far has centered around US diplomatic activity, the practice of exchanging visits for concessions is by no means a uniquely American phenomenon. For

[2] Quote from an interview conducted by Myrick and Weinstein (2020).

[3] https://history.state.gov/historicaldocuments/frus1964-68v29p1/d354

[4] https://history.state.gov/historicaldocuments/frus1955-57v14/d175

instance, we see similar tactics employed by both British Prime Minister Tony Blair and French President Nicolas Sarkozy in their dealings with Libyan leader Muammar Gaddafi in 2007. Both European leaders leveraged a high-profile diplomatic visit in exchange for advantageous commercial deals for their own domestic firms (The New York Times 2007; BBC News 2007). These were the publicly known considerations in the deals; the allegation subsequently emerged that Gaddafi illicitly paid €50 million to Sarkozy's 2007 election campaign in advance of the visit, a crime for which Sarkozy has since been charged by French prosecutors (McAuley 2018).

What does existing literature tell us about these kinds of international exchanges? In general terms, we are interested in a situation in which two actors are bargaining in front of a third party who draws inferences from the bargaining outcome that she observes. Interactions of this sort have been examined in previous formal literature, with a particular focus on the incentives for "posturing" that arise, and the consequences of transparency and information asymmetry for bargaining outcomes (Groseclose and McCarty 2001; Stasavage 2004; Perlroth 2019). Our situation is unique, however, in that the object of the bargaining is not a division of resources or a policy output; rather, the object of the bargain is *information*, in the form of a costly, "money-burning" signal sent to a third party. In this sense, the puzzle we present is similar to that considered by Vreeland and Dreher (2014) in the context of vote-buying in the UN Security Council, where, as the authors describe, "the central political commodity that is bought and sold is legitimacy". While Vreeland and Dreher's study is primarily an empirical investigation of aid-for-voting exchanges, we provide a formal theoretical analysis which highlights the tensions and complexities inherent to the concessions-for-signalling exchanges under examination.

More directly related to our substantive context of analysis is a newly emerging body of quantitative and formal literature on public diplomacy and diplomatic visits (Malis and Smith 2019, 2021; McManus 2018).[5] These studies generally theorize diplomatic visits as public signals of support from one leader to another.[6] McManus (2018) conceptualizes visits as tied-hands deterrent signals in the face of threats from foreign adversaries. Malis and Smith (2019, 2021) examine how visits deter domestic challenges against the recipient

[5]Complementing this body of work on the public-facing aspects of in-person diplomacy is a set of studies focusing on its private aspects: see, for instance, Holmes and Yarhi-Milo (2016).

[6]Matush (2020), in contrast, considers how public diplomacy can be used for foreign antagonism rather than support.

leader. The present analysis differs substantially from these previous studies in two key respects: first, the model endogenizes the size of the quid pro quo through a bargaining process; and second, the domestic leader is treated as a strategic actor. An implication of these two innovations is that the signal receiver in our model knows that a visit was "purchased", and thus conditions their inference on whatever price was paid. This in turn alters the bargaining incentives of both the domestic and foreign leader. As a consequence, our model reveals that the kind of diplomatic exchanges theorized by Malis and Smith (2019, 2021) are not generally incentive-compatible for both parties. We instead show that the conditions under which mutually beneficial quid pro quo diplomacy can occur are precisely the conditions under which it would be least expected, as we elaborate in Section 6 below. Further, our analysis uncovers the crucial role of transparency in diplomatic exchange: while the visit itself must be public and visible to third parties, the negotiations surrounding the visit must be conducted out of public view—and the degree of transparency has important distributional implications for bargaining outcomes.

2 A Model of Concessions, Visits, and Political Survival

Our model features three players: a domestic incumbent, L; her political challenger, C; and a foreign leader, F. The model has three phases: 1) bargaining; 2) visiting; and 3) domestic political competition. In the bargaining phase, the salience (S) of a concession arises stochastically, and L and F bargain (under varying bargaining protocols and informational structures) over the size of the concession, z, that L will provide in exchange for a visit. F then decides whether or not to conduct the visit and reap the concession. Following the bargaining and the occurrence (or not) of a visit, domestic competition occurs, in which C decides whether or not to attempt to remove L from power. We reduce domestic competition to two dimensions: the cost K that C incurs for mounting a challenge against the leader, and the probability θ that the leader can survive a challenge. The challenger fully learns the first factor, while the foreign leader sees a noisy signal of the latter.

2.1 Bargaining

The game begins with nature stochastically determining whether a salient opportunity exists for L to provide F a concession. We represent this as a random variable $S \in \{0,1\}$. The realization of S is observed by both L and F. The domestic challenger C has a prior belief that $Pr(S=1) = \sigma \in (0,1)$, and that belief is updated over the course of the game, as discussed below. We refer to σ as "expected salience", that is, C's expectation that F wants a favor from L.

Absent such a salient opportunity, the bargaining phase concludes trivially with no concession agreed upon, and the game moves directly to the political competition phase. If a salient opportunity arises, then bargaining commences and L and F negotiate the size of the concession that L is to provide F in exchange for a visit. We assume a simple take-it-or-leave-it bargaining framework: one leader, either L or F, proposes a concession z to be proffered upon the completion of a visit, and the other leader either accepts or rejects the deal.

We vary the bargaining context along two dimensions: 1) which leader has proposal power, and 2) whether the negotiations are "open" or "closed".

First, with respect to proposal power: If L is the proposer, then she offers a concession of size $z \geq 0$ in return for a visit, which F then either accepts or rejects. If F is the proposer, then the roles are reversed: F makes a demand of z, which L either accepts or rejects. We note that accepting an agreement at this stage does not entail a commitment on F's part to conduct a visit; rather, it is a commitment on L's part to provide a concession of size z in the event that F does decide to visit (which we discuss in the next phase).

Second, considering the transparency of negotiations: Visits are public events seen by all. However, we vary the extent to which C observes other aspects of the game. Under the "open" bargaining protocol, the first scenario which we examine, C observes all aspects of the negotiations: that is, C observes the realized salience S, the concession size z that was offered or demanded, and whether the deal was accepted or rejected. In the "closed" protocol, which we believe more closely resembles the reality of diplomatic negotiations, C does not observe concession salience or any aspect of the pre-visit bargaining that occurs between F and L. Within the closed setting we vary two features of the model setup: first, we vary the expected salience, σ; and second, we vary the probability $q \in [0,1]$ that, in the event that a visit occurs, the concession value z is revealed to C. We refer to this probability q as "concession transparency". If C does not observe z, he still infers that a concession was made, but must form a conjecture as to the size of the concession.[7]

2.2 Visits

Once a deal z is agreed to, F must decide whether to carry out the visit and collect the concession. From F's perspective, the decision to visit entails balancing costs and benefits, which are conditional on L's survival in office. Before embarking on the visit, F sees a noisy signal A of L's strength θ, which captures the likelihood that L can survive a challenge from her domestic rival (described below in detail). For clarity of exposition, we specify the game such that F observes this signal after the negotiations are complete; however, as we discuss in the appendix, the results are robust to alternative assumptions about the timing of F's information acquisition.

If F conducts a visit, he incurs a fixed cost τ: this includes the opportunity cost of F's time, as well as any associated transport, security, and administrative costs. F pays this cost regardless of whether or not L survives in office. In the event that L is removed from office following the visit, F pays an additional cost of ρ; we refer to this as a reputational cost, following Malis and Smith (2021). Weighed against these costs is the benefit F enjoys from the concession z. If L remains in office, this concession is received with certainty; if L is removed, F retains the concession with probability $r \in [0, 1]$. Thus r captures the immediacy with which the concession can be delivered following the negotiation, or the difficulty that L's successor would face in revoking the concession.

2.3 Domestic political competition

In the final stage of the game, the challenger, C, may attempt to remove the leader. For expositional purposes we focus on domestic political competition, though the model readily accommodates a foreign challenger. Many factors shape political contestation. We focus here on two dimensions: θ and K, which we refer to as the leader's "strength" and the challenger's "cost" respectively. θ represents the probability that L can withstand a

[7]While the "open" versus "closed" distinction bundles together multiple features of the bargaining structure, any intermediate cases are effectively nested within the two extreme cases. See Section 8.3 of the appendix.

removal attempt by C, while K represents C's cost for making such an attempt. All actors have common priors about θ and K. The challenger, C, fully learns his cost K before deciding whether or not to challenge. We assume that F learns nothing about the cost dimension, but does see a noisy signal about L's strength, θ. Thus while C maintains an informational advantage over F with respect to one dimension of political contestation, he can still learn something about the second dimension from observing F's decision to visit or not.

C's payoff from successfully removing L is 1 and we normalize C's payoff from L remaining in office to 0. Given regime strength θ, C's expected payoff from attempting to remove L is

$$(1-\theta)1-K$$

We assume K is uniformly distributed on the internal $[0,1]$.

The domestic leader L receives an office-holding benefit of Ψ if she retains power, while her payoff from being deposed is normalized to 0. If a visit occurs, she incurs the cost z for making the concession.

2.4 Signals and beliefs

Foreign leader F observes a noisy signal of L's strength θ. This signal A takes values between 0 and n. For intuition (although we do not restrict A to integers) we can think of A as the number of "heads" from n biased coin flips, where the probability of heads in each trial is θ. By this interpretation, each coin flip would represent a new piece of evidence uncovered about L's strength in office, and heads would denote evidence indicating that L is strong, and hence capable of surviving in office to deliver a concession. The number of trials, n, thus provides a convenient metric for the precision of F's signal.

All players share a common prior belief of θ, which lies between 0 and 1. We specify this prior distribution, $G(x) = Pr(\theta < x)$, to be the Beta distribution (with parameters α and β), which is a flexible distribution on the domain [0,1]. The associated probability density is $g(\theta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}$, where $\Gamma(z) = \int_0^\infty x^{z-1}e^{-x}dx$.

We exploit two useful features of the Beta distribution (as have others, for instance Alt, Calvert and Humes (1988)). First, the expected value of θ is $\frac{\alpha}{\alpha+\beta}$. Second, the Beta distribution has a simple Bayesian update. Given regime strength θ, the probability density with which $A = x$ is

$$p(x|\theta) = \frac{\Gamma(n+1)}{\Gamma(x+1)\Gamma(n-x+1)}\theta^x(1-\theta)^{(n-x)},$$

the likelihood of x heads from n coin flips where the coin is weighted to come up heads with probability θ (for integers the terms involving the Gamma functions simply reduce to the binomial coefficients). Let $p(x) = \int_0^1 p(x|\theta)g(\theta)d\theta$ and $P(a) = \int_0^a p(x)dx$. Therefore, $Pr(A < a) = \frac{P(a)}{P(n)}$ and $Pr(A \geq a) = \frac{P(n)-P(a)}{P(n)}$.[8]

Given the signal A, Bayesian updating yields the posterior belief that θ is Beta distributed with parameters $\alpha + A$ and $\beta + (n - A)$, so the updated expectation that L can survive a challenge is $E[\theta|A] = \frac{\alpha+A}{\alpha+\beta+n}$.

2.5 Summary of model setup

To summarize, the sequence of the game form is as follows:

1. Nature stochastically determines whether a salient opportunity for a favor arises, $S \in \{0, 1\}$, with $Pr(S = 1) = \sigma$.
2. If $S = 1$, then bargaining occurs. Depending on the bargaining protocol, either L offers concession z, which F can accept or reject, or F demands concession z which L can accept or reject.
3. F sees signal A and decides whether to visit L.
4. Under open bargaining, C observes all aspects of the negotiations. Under closed bargaining, C observes whether a visit occurred and, should a visit occur, C observes the concession granted with probability q.
5. Domestic political competition: C either challenges or abstains. If C plays challenge, then L is removed with probability $1 - \theta$.

Payoffs and notation are in Tables 1 and 2. We characterize Perfect Bayesian Equilibria under different bargaining protocols.

[8] The model considers $A \in [0, n]$. The standard binominal setup considers only integers. While $\sum_{x=0}^{n} p(x) = 1$, the integral $\int_0^n p(x) \neq 1$. Hence throughout we need to standardize by $P(n)$.

Table 1: Payoffs

F, Foreign Power	L Survives	Regime Change
Visit	$Sz - \tau$	$S\,r\,z - \tau - \rho$
Non-visit	0	0

L, Leader	L Survives	Regime Change
Visit	$\Psi - z$	$-z$
Non-visit	Ψ	0

C, Challenger	Status Quo	Regime Change
Challenge	$-K$	$1 - K$
Abstain	0	n.a.

Table 2: Notation

$K \sim U[0,1]$	C's cost of challenging
$\theta \in [0,1]$	L's regime strength, with prior distribution $Beta(\alpha, \beta)$
$A \in [0,n]$	F's private signal of regime strength
$z \geq 0$	Concession offered in exchange for a visit
$\Psi > 0$	L's valuation of holding office
$S \in \{0,1\}$	Salience of concession for F, with prior $Pr(S=1) = \sigma$
$r \in [0,1]$	Probability that F retains a concession following L's removal
$\tau > 0$	F's fixed/opportunity cost for visiting
$\rho > 0$	F's conditional/reputational cost for visiting if L is subsequently removed
$\sigma \in (0,1)$	Expected salience: C's prior belief of $Pr(S=1)$
$q \in [0,1]$	Concession transparency: prob. C observes z after visit under closed bargaining

3 Visit Subgame

We begin with a general characterization of the subgame that follows the negotiation stage, which applies to both open and closed bargaining settings.

3.1 F's incentive to visit

The core incentive is that F wants to visit strong leaders. First note that absent a salient opportunity for a concession ($S = 0$), no bargaining occurs, and the game proceeds directly to domestic political competition. Thus our analysis of F's visit decision focuses on the case of $S = 1$.

Suppose the negotiated concession is z, and F sees signal A, and following a visit the challenger will attempt removal with probability κ_v. Then F's expected payoff from a visit is

$$\begin{aligned} V(A,z) &= Pr(L \text{ survives}|visit, A)(z-\tau) + Pr(L \text{ deposed}|visit, A)(rz-\tau-\rho) \\ &= z - \tau - \underbrace{\kappa_v}_{\text{Pr(challenge)}} \underbrace{E[1-\theta|A]}_{\text{Pr(deposed|challenge)}} ((1-r)z+\rho) \end{aligned} \tag{1}$$

If C attempts removal, then L is deposed with probability $1-\theta$. The likelihood of a high signal increases in L's strength θ. Therefore F is more likely to visit as the signal A increases: $E[1-\theta|A] = \frac{n-A+\beta}{\alpha+\beta+n}$, for $A \in [0,n]$. Equation (1) is strictly increasing in z and A. If he refrains from visiting, the foreign leader's payoff is normalized to 0, so F only visits when $V(A,z) \geq 0$. Thus there is a unique concession size that makes F indifferent between visiting and not.

Let the strictly decreasing function $z(a)$ characterize the indifference concession associated with each signal a. Conversely, define the inverse function, $a(z)$, as the signal that makes F indifferent between visiting and not given concession z. Further, define

$$z_0 = \frac{\frac{\beta\rho(\beta+n)}{(\alpha+\beta)(\alpha+\beta+n)} + \tau}{1 - \frac{\beta(1-r)(\beta+n)}{(\alpha+\beta)(\alpha+\beta+n)}} \qquad \text{and} \qquad z_n = \frac{\frac{\beta^2\rho}{(\alpha+\beta+n)^2} + \tau}{1 - \frac{\beta^2(1-r)}{(\alpha+\beta+n)^2}}$$

as the extreme concessions that make F indifferent between visiting following the weakest and strongest possible messages respectively: $V(A=0,z_0)=0$ and $V(A=n,z_n)=0$. Note that $z_n < z_0$: a large concession ($z=z_0$) is necessary to make F indifferent given the lowest signal ($a=0$), and a small concession ($z=z_n$) is necessary to make F indifferent given the highest signal ($a=n$).

To summarize:

Lemma 1 *For concessions $z \in (z_0, z_n)$, there is a decreasing monotonic function $a(z)$ such that $V(a(z),z)=0$. F visits if and only if he sees a signal $A \geq a(z)$.*

3.2 Visits deter domestic political challenges

In the final move of the game, the challenger C decides whether to challenge or abstain. L survives a removal attempt with probability θ. Given information I, C only challenges when the expectations of success justify the cost of attempting removal:

$$\underbrace{E[1-\theta|I]}_{Pr(success)} - \underbrace{K}_{\text{Cost of attempt}} \geq 0$$

There is critical threshold $k = E[1-\theta|I]$ such that C attempts removal when $K \leq k$. The critical threshold depends upon the information revealed by bargaining and visits. To simplify notation, we use k_0 to denote C's challenge threshold given that there was no opportunity for a visit ($S = 0$ in open bargaining, or $\sigma \to 0$ in closed bargaining). The thresholds $k_v(z)$ and $k_{nv}(z)$ indicate C's critical cost given the occurrence and absence of a visit, respectively, given C's knowledge (in open bargaining) or conjecture (in closed bargaining) that the agreed-upon concession was of value z.

If C knows that there was no salient opportunity for a visit ($S = 0$), then the absence of a visit is uninformative, and C cannot update beyond his prior:

$$k_0 = E[1-\theta] = \frac{\beta}{\alpha+\beta} \quad \text{and} \quad Pr(\text{challenge}) = Pr(K < k_0) = k_0 = E[1-\theta]$$

Visits affect domestic political competition by altering C's beliefs about regime strength. The occurrence of a visit reveals that F saw a signal $A \geq a(z)$. Given this information, C's expectation of the likelihood of a successful challenge declines, since $E[1-\theta|A \geq a(z)] \geq E[1-\theta]$. In contrast, if F forgoes the opportunity to visit (but C knows there was an opportunity), then C infers that L is weaker than initially thought: $E[1-\theta|\text{non-visit}] = E[1-\theta|A \leq a(z)] \leq E[1-\theta]$.

We can now state equilibrium behavior in the visit subgame:

Proposition 1 ***Informative equilibrium:*** *Suppose the bargaining phase results in a concession of z that is known to all players. If $z \in (z_n, z_0)$, then F visits if and only if $A \geq a(z)$ where $a(z)$ is the implicit function that that solves*

$$V(a,z) \quad = \quad z - \tau - E[1-\theta|A=a]E[1-\theta|A \geq a]((1-r)z+\rho) = 0 \tag{2}$$

Following a visit C attempts removal if and only if $K \leq k_v(z) = E[1-\theta|A \geq a(z)]$, and following non-visit, C attempts removal if and only if $K \leq k_{nv}(z) = E[1-\theta|A < a(z)]$, with $k_v(z) < k_0 < k_n v(z)$.

Consistent with standard refinements (Banks 1991), we impose the following restriction on out-of-equilibrium beliefs:

Assumption 1 *If $a(z) \geq n$ (such that no visits occur), then should a visit occur let $E[1-\theta|visit] = \frac{\beta}{\alpha+\beta+n}$. If $a \leq 0$ (such that visits always occur), then let $E[1-\theta|non\text{-}visit] = \frac{\beta+n}{\alpha+\beta+n}$.*

Intuitively, if C were to observe a visit when a visit should never have occurred, he draws the most positive inference of L's strength; and conversely, if C were to observe a non-visit when a visit should necessarily have occurred, he draws the most negative inference of L's strength. With this restriction in place:

Proposition 2 ***Pooling equilibria:*** *If the concession is small, $z \leq z_n$, then F never visits and following no visit C attempts removal if $K \leq k_0$. If the concession is large, $z \geq z_0$, then F always visits and following a visit C attempts removal if $K \leq k_0$.*

Note that Equation (2) is simply Equation (1) with the substitution that $\kappa_v = k_v = E[1-\theta|A \geq a]$. For any given z, the equilibrium to the visit subgame is unique.

Corollary 1 *For an fixed concession z, visits become more likely as r increases and as τ and ρ decrease: $\frac{\partial a(z)}{dr} \leq 0$, $\frac{\partial a(z)}{d\tau} \geq 0$ and $\frac{\partial a(z)}{d\rho} \geq 0$. The inequalities are strict for $z \in (z_n, z_0)$.*

Holding fixed the size of the concession, for a greater risk of default $(1-r)$ or greater costs of visiting (τ and ρ), F requires a more favorable signal in order to justify making a visit.

Corollary 2 *For the special case $\alpha = \beta = 1$,*

$$a(z) = \frac{3}{2} + n - \frac{1}{2}\sqrt{\frac{8(n+2)^2(z-\tau) + \rho + (1-r)z}{\rho + (1-r)z}} \tag{3}$$

for $z \in (z_n, z_0)$

Alternatively, stated in terms of concessions,

$$z(a) = \frac{\rho(a-n-2)(a-n-1) + 2(n+2)^2\tau}{a^2(r-1) - a(2n+3)(r-1) + (n+2)(nr+n+r+3)} \tag{4}$$

Visits discourage removal attempts ($k_v(z) < k_0$) and non-visits increase the likelihood of removal attempts ($k_{nv} > k_0$), for $z \in (z_n, z_0)$. However, the likelihood of a visit and the domestic political consequences of a visit depend upon the size of the concession on offer.

Proposition 3 ***Effect of Concessions on Occurrence and Impact of Visits**: As concessions increase,*

1. *visits become more likely:* $\frac{dPr(visit)}{dz} \geq 0$
2. *the perception of L's strength following a visit decreases:* $\frac{dE[\theta|visit]}{dz} \leq 0$
3. *the perception of L's strength following no visit decreases:* $\frac{dE[\theta|non\text{-}visit]}{dz} \leq 0$

These inequalities are strict if $z \in (z_n, z_0)$.

The intuition is straightforward: visits signal regime strength and enhance L's survival. Hence L wants to encourage F to visit. However, Proposition 3 shows that making visits more attractive undermines their domestic political value. In the extreme, if C knows that L made a concession of size z_0 (which is large enough to induce F to visit even for the worst possible signal), then the visit fails to signal L's strength. As the concession size decreases, F becomes less likely to visit—but the deterrent value of a visit, should it occur, grows stronger. The tension between wanting to increase the likelihood of a visit and undermining the domestic political value of a visit shapes equilibrium bargaining behavior.

4 Open Bargaining

Under open negotiations, C, observes all details of the bargaining. In this case, we have the following general result:

Proposition 4 *Under open bargaining, any concession* $z \in (z_n, z_0)$ *that induces threshold strategy* $a(z) \in (0, n)$ *increases L's risk of deposition relative to no agreement (*$z < z_n$*).*

When the leaders negotiate in the open, the lottery between a survival-improving visit and a deposition-enhancing non-visit reduces L's aggregate survival prospects relative no information being communicated. This result holds regardless of whether F or L has proposal power.

Figure 1: Probability of Regime Change (RC) given Threshold $a(z)$ under Open Bargaining

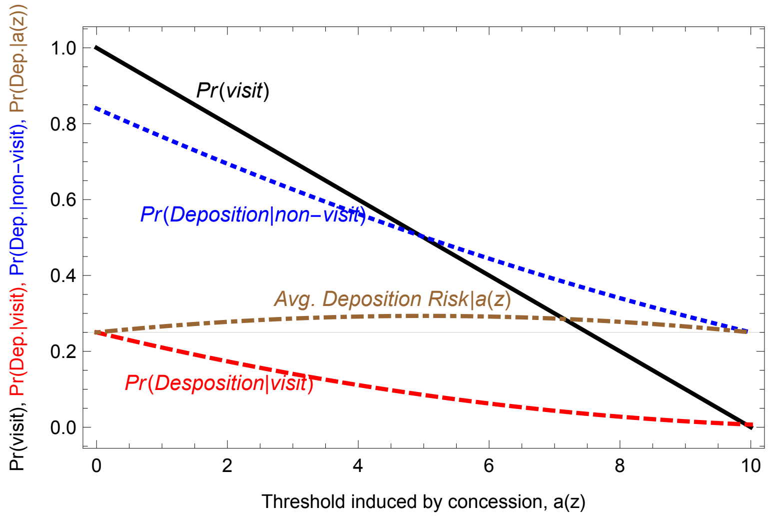

Figure 1 illustrates Proposition 4 for the special case of a uniform prior ($\alpha = \beta = 1$). The horizontal axis shows the threshold $a(z)$ induced by concession z. As z increases, $a(z)$ decreases, and F visits for a greater range of signals. The probability that F visits is shown by the solid line. The figure also shows the probability that L is deposed following a visit (dashed, $Pr(Deposition|visit)$) and no visit (dotted, $Pr(Deposition|\text{non-visit})$). If a bargain is reached ($z \in (z_n, z_0)$), then L's expected risk of being deposed is:

$$\begin{aligned}\text{Avg. Deposition Risk}|a(z) &= Pr(visit)Pr(Deposition|a(z)|visit) \\ &+ (1 - Pr(visit))Pr(Deposition|a(z)|\text{non-visit}) \qquad (5)\end{aligned}$$

This aggregate deposition risk, shown by the dot-dashed hump-shaped curve, is higher than if no additional information were revealed. Thus it is clear to see that L does not want any deal that could result in a visit: doing so would entail paying a direct cost z in order to decrease her own expected survival prospects. Instead L would propose $z < z_n$ and reject any deal offered by F in which $z \geq z_n$.

Proposition 5 *Under open bargaining, concessions are never agreed and quid pro quo visits never occur. If L is proposer, then she proposes $z < z_n$. If F is proposer, then L rejects any offer $z \geq z_n$. C does not revise his beliefs of θ after observing the absence of a visit.*

The potential for quid pro quo diplomacy carries only downside risk for L, such that she would prefer to commit to a strategy of never offering a concession in exchange for a visit. Given the fully transparent bargaining protocol, such a commitment strategy is credible: C observes that no concession is on offer, and thus does not update negatively upon observing the absence of a visit. As we will see in the next section, removing transparency in negotiations has the effect of removing L's ability to credibly commit to this no-concession strategy, forcing her into deals which leave her worse off on average.

5 Closed Bargaining

While the open bargain setting provides a useful analytical benchmark, it is substantively unrealistic to assume that domestic actors outside of government have the opportunity to observe all details of diplomatic negotiations. Indeed, observers may not even know that negotiations are occurring. In the closed setting C does not observe the salience of concession opportunities or the bargaining over concessions. Within this closed setting, we vary expected salience, σ, and concession opacity, q.

Recall that C holds a prior belief that $S = 1$ with probability σ, and recall that no visits occur when $S = 0$. Upon observing the absence of a visit, C infers that one of two things happened: either $S = 0$, so there was no opportunity for a visit; or $S = 1$, but F did not observe a strong enough signal to warrant a visit, given the concession z that was on offer. Suppose that in the latter case C conjectures that the size of the negotiated concession was w. Then his posterior belief of S is given by Bayes' Rule:

$$Pr(S = 1|\text{non-visit}, w) = \frac{P(a(w))\sigma}{P(a(w))\sigma + 1 - \sigma}$$

Absent a visit, C's expectation of regime weakness is

$$\begin{aligned} E_C[1-\theta|\text{non-visit}, w] &= Pr(S = 1|\text{non-visit})E[1-\theta|A < a(w))] + (1 - Pr(S = 1|\text{non-visit}))\frac{\beta}{\alpha+\beta} \\ &= \frac{P(a(w))\sigma E[1-\theta|A < a(w))] + (1-\sigma)P(n)\frac{\beta}{\alpha+\beta}}{P(a(w))\sigma + (1-\sigma)P(n)} \end{aligned}$$

Observe that this value is increasing in σ: when it is ex-ante more likely that there exists a salient opportunity for a diplomatic exchange, the absence of a visit becomes a more damaging signal ex-post.

5.1 F as proposer

Closed-door bargaining outcomes vary depending on which leader, F or L, has proposal power. First let us consider the case where F makes a demand z which L can accept or reject.

Proposition 6 *If $q \in [0,1]$, $\sigma \in (0,1)$ and F is proposer, then on the equilibrium path, when $S = 1$, F proposes z such that*

$$\frac{z}{\Psi} = \frac{\frac{\beta}{\alpha+\beta}P(n) - E[1-\theta|A < a(z))]P(a(z))}{P(n) - P(a(z))}E_C[1-\theta|non\text{-}visit, z] - E[1-\theta|A \geq a(z))]^2, \tag{6}$$

which L accepts. The concession z is increasing in σ. As office holding becomes the dominant motive ($\Psi \to \infty$), $z \to z_0$.

It is instructive to compare this result against the open-door bargaining result. In the open case, suppose F were to demand some concession $z \in (z_n, z_0)$. As demonstrated above, L prefers that F's action be uninformative, and also prefers not to pay the direct cost of the concession. So she turns down the deal, and C knows that the absence of a visit is uninformative.

Such behavior, however, cannot be supported in equilibrium under closed bargaining. Suppose that C conjectured that L's strategy was to turn down any offer. Off the equilibrium path, if C were to observe a visit, he would have to infer (by Assumption 1) that F saw a favorable signal of L's strength, and would update his beliefs of θ accordingly. Then L has a clear incentive to deviate and accept the lowest-z deals that F would offer. This leads C to update more negatively upon observing a non-visit, which in turn drives up the price L is willing to pay. In equilibrium, L ends up proffering a substantial concession in order to secure an average deposition risk that is at least as large as what she obtained for free under open bargaining. The concession size that L is willing to pay is increasing in the expected salience σ, because C's inference of regime strength (and thus L's survival prospects) when no visit occurs is decreasing in σ.[9]

We characterize this result as a diplomatic resource curse. By virtue of being perceived as having something to offer the foreign power (that is, σ being large), the domestic leader

[9] The results in Proposition 6 hold for $\sigma \to 1$. However, if C is completely certain, $\sigma = 1$, then F can shake down L to an even greater extent. We characterize this pathological case in the appendix (Proposition 8).

faces domestic costs for refraining from diplomatic engagement. She cannot credibly commit to having *not* offered a concession, and so she faces a strong incentive to make a large concession that is likely to secure a diplomatic visit. If instead it was thought to be unlikely that L had a salient concession to offer (that is, if σ were small), then in the off chance that a salient opportunity did arise, L would come to the diplomatic table with greater bargaining leverage due to her more favorable reservation value for turning down a deal. Though not explicitly allowed in our model, this logic suggests an incentive for some leaders either to isolate themselves diplomatically—making clear to their domestic audiences that foreign powers are not demanding any sort of concession from them—or to make their negotiations more transparent, so that audiences cannot suspect them of having offered a larger concession than they did.

5.2 L as proposer

Now we turn attention to the case of closed-door bargaining with L as proposer. In the appendix we characterize the equilibria in a more general setting. For ease of presentation, here we focus on the case where office-holding is the dominant motive.

Proposition 7 *If L is proposer and office holding incentives dominate ($\Psi \to \infty$), then on the equilibrium path L offers z that solves:*

$$E[1-\theta|A \geq a(z)]\big(E[1-\theta|A = a(z)](1+q) - E[1-\theta|A \geq a(z)]q\big) \\ - \frac{E[1-\theta|A < a(z)](P(n)\frac{\beta}{\alpha+\beta}(1-\sigma) + P(a(z))E[1-\theta|A < a(z)]\sigma)}{P(n)(1-\sigma) + P(a(z))\sigma} = 0 \qquad (7)$$

and visits occur if $A \geq a(z)$.

L trades off the likelihood of receiving a visit against the impact of a visit, should it occur. If L offers a small concession, then F rarely visits—but, if a visit occurs, the visit greatly increases L's survival prospects. If L offers a large concession, then visits are common, but their impact on domestic politics is small. When officeholding motives dominate, L finds the threshold $a(z)$ that maximizes expected survival and then offers the concession that induces that threshold. The optimal threshold $a(z)$ represents a compromise between the probability of a visit and domestic political benefits of a visit.

Figure 2: L's Choice of Concessions Behind Closed Doors

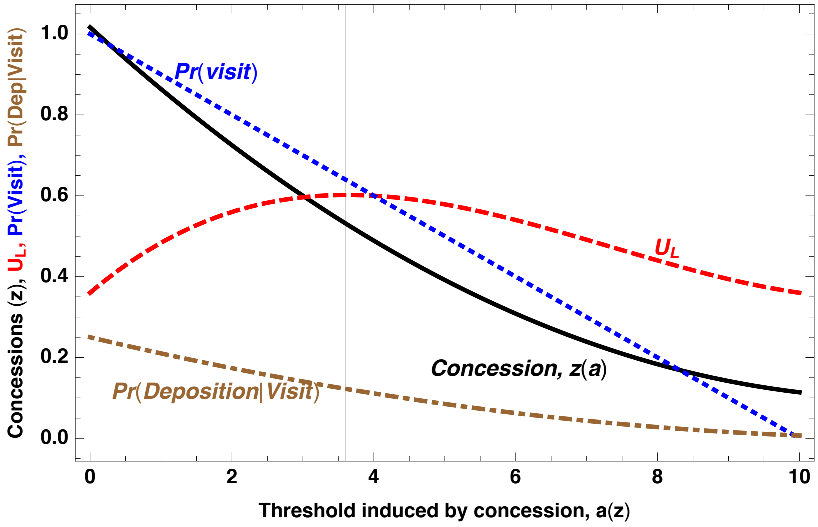

To illustrate the tradeoff between likelihood of a visit and a visit's impact, we return to the special case of the uniform prior, $\alpha = \beta = 1$, with $\sigma \to 0$, $q = 1$, and dominant office-holding incentives ($\Psi \to \infty$). This case is depicted visually in Figure 2.[10] The threshold $a(z)$ is the horizontal axis. The solid decreasing line shows the concession z necessary to induce each threshold a. The downward sloping dotted line shows $Pr(visit)$. The lower dash-dot line shows the probability of regime change given the occurrence of a visit. The dashed line shows L's (rescaled) expected payoff for the each induced threshold.

A large concession induces a low threshold (LHS of figure, $a(z)$ close to 0), and so visits become very likely. However, such visits are of little value to L, as they provide only a weak signal of strength. At the opposite extreme, if L offers only a small concession, then the induced threshold is high, meaning visits are unlikely, but visits provide a strong signal of strength when they do occur. L's payoff is maximized by a intermediate-sized concession, such that visits occur reasonably often and provide a moderate signal of strength.

We can examine further how bargaining outcomes depend on concession transparency and expected salience:

Corollary 3 *If L is proposer: As concession transparency q increases, the concession size z becomes smaller and visits become less likely. As expected salience σ increases, the concession z increases and visits become more likely. When office holding is the dominant*

[10]The figure was constructed assuming $n = 10$, $\Psi \to \infty$ (and L's payoff rescaled for the graph), $r = 1$, $\rho = 2$ and $\tau = 0.1$

concern ($\Psi \to \infty$)*, if either* $q \to 0$ *or* $\sigma \to 1$*, then* $z \to z_0$*.*

Expected salience increases the size of concession that L offers, following the resource-curse logic outlined above, and thus makes visits more likely. As C believes it to be increasingly likely that F wants a favor from L, the lack of a visit sends an increasingly negative signal of L's strength: $\frac{dE_c[1-\theta|\text{non-visit},z]}{d\sigma} > 0$. To avoid this negative signal, L offers a large concession that makes F highly likely to visit.

The domestic leader's equilibrium offer is also determined by concession transparency: the size of the concession is decreasing in the probability that it gets revealed to C upon delivery. To see why, first note that with full concession transparency, $q = 1$, L could set the concession size to optimally balance the likelihood of a visit against the informativeness of a visit; if a visit occurs, C observes the associated concession and draws exactly the inference that L intended. But now consider what happens in the other extreme, $q = 0$, so that when a visit occurs, C does not observe the concession that was made. If C conjectures that L has offered the same concession as she had in the $q = 1$ case, then L faces the incentive to increase the offer, because doing so increases the probability of a visit without undermining the signaling value of the visit. Thus L will deviate to a larger concession; but, knowing that L faces this incentive, C increases his conjecture of the offered concession in turn. Concession transparency limits the size the concession, because without transparency L cannot resist the temptation to bid up her offer.

As a final consideration, with the uniform prior, we can provide an explicit statement of the equilibrium concession size and corresponding threshold, and examine the role of F's intelligence:

Corollary 4 *For* $\alpha = \beta = 1$*, low expected salience* ($\sigma \to 0$) *and dominant office incentives* ($\Psi \to \infty$)*,* L *offers concession*

$$z = \frac{\rho\left(\sqrt{n(n+2)+(q+1)^2}+n+1\right)\left(\sqrt{n(n+2)+(q+1)^2}+n+q+3\right)+2(n+2)^2(q+2)^2\tau}{(r-1)\left(\sqrt{n(n+2)+(q+1)^2}+n+1\right)\left(\sqrt{n(n+2)+(q+1)^2}+n+q+3\right)+2(n+2)^2(q+2)^2} \tag{8}$$

that induces

$$a(z) = \frac{(n+1)(1+q)}{2+q} + \frac{\sqrt{n^2+2n+(1+q)^2}}{2+q} \tag{9}$$

As n *increases, visits become more likely* $\left(\frac{d\frac{a(z)}{n}}{dn} < 0\right)$.

L and F's payoffs are both increasing in n. The quid pro quo of buying visits is more attractive to both parties as the foreign power has access to better intelligence.

6 The Price of the Quid Pro Quo

Given the results of Propositions 6 and 7, we can now address the question posed at the outset of the paper: what is the price of a diplomatic visit? The preceding analysis produced a number of comparative statics on bargaining outcomes as a function of various features of the bargaining protocol and attributes of the two leaders, which we summarize here.

Transparency vs. Opacity When C is fully aware of salient opportunities and the details of negotiations (as in our "open door" bargaining protocol), then an exchange of concessions for visits is never in the recipient leader's interest. Quid pro quo diplomacy requires that bargaining occurs behind closed doors. However, our analysis reveals a more nuanced effect of transparency in diplomacy: while the domestic leader benefits from opacity in negotiations, she prefers transparency in the actual delivery of the negotiated concession. When neither the negotiations nor the concessions themselves are observed by the domestic challenger, the size of the concession increases, and visits become more frequent but less informative.

Bargaining power: If L makes the proposal, she offers a moderate concession. Her moderate concession balances a desire to increase the likelihood of a visit while maintaining its deterrent value. As bargaining power shifts from L to F, bargaining outcomes shift to favoring F over L. The size of concessions increases, visits become more frequent, and the visits that occur have a decreased impact on L's survival.

Expected Salience: When F is likely to want a favor, L suffers from a diplomatic resource curse. Because C anticipates that L has something of value to offer F, the fact that F does not make a visit in order to obtain the concession sends a strong negative signal of L's strength. To avoid this signal of weakness, L pays more for a visit, visits become more likely and less effective. In contrast, when it is seen as unlikely that F wants a favor, the absence of a visits is less politically detrimental to L; this drives the price of the visit down, which in turn implies that any visits that do occur provide an especially powerful signal of L's strength.

F's cost: As F's costs of visiting increase, so does the price needed to offset those costs. There are three distinct costs that F must factor into his visit decisions: the material and opportunity costs of travel, τ; the reputational cost of associating with a soon-deposed leader, ρ; and the risk that any agreed-upon deal does not get implemented, $1 - r$. A stop on a regional tour or a pull-aside at a multilateral summit can be purchased more cheaply than a trip undertaken solely for the bilateral visit. Dictators and human rights abusers will likely see a steeper price for their diplomatic engagements. A policy concession requiring long-term implementation will need to be greater than its cash equivalent, as the policy concession carries the risk that L fails to survive in office long enough to implement it.

Ex-ante survival prospects: Prior expectations of the likelihood that L will survive affect the size of concessions and the likelihood of visits. Intuitively, ex-ante survival prospects affect both the supply and demand of visits. If L is likely to be deposed, then F needs larger concessions to compensate for the risk he assumes, due to the potential reputational cost as well as the possibility that the concession does not get implemented. When pessimistic about her survival prospects, L values visits highly and is willing to make larger concessions. Thus when L's regime is perceived to be unstable, the supply of visits is low and the demand for them high, so the price is driven upwards. Figure 3 illustrates the impact of survival prospects on concessions when L has proposal power and officeholding incentives are dominant ($\Psi \to \infty$).

The horizontal axis in Figure 3 is $E[\theta] = \frac{\alpha}{\alpha+\beta}$. The dashed line shows L's optimal concession and the solid line shows the threshold $a(z)$ that the concession induces.[11] When L is anticipated to be strong, she offers relatively small concessions and F only visits if he sees a relatively high signal. The relative rarity of visits means that a visit is a powerful signal of strength. In contrast, when L is perceived to be weak, L makes more generous concessions that induce a lower threshold $a(z)$, meaning that F visits for a wider range of private signals. L is willing to spend more for a weaker signal of strength when she is perceived to be weak because her precarious situations makes any signal of strength valuable. Of course, the increased concession does not necessarily mean visits are more

[11]The figure is constructed using $n = 10$, $q = 1$ and $\alpha + \beta = 6$.

Figure 3: L's Anticipated Strength and Concessions

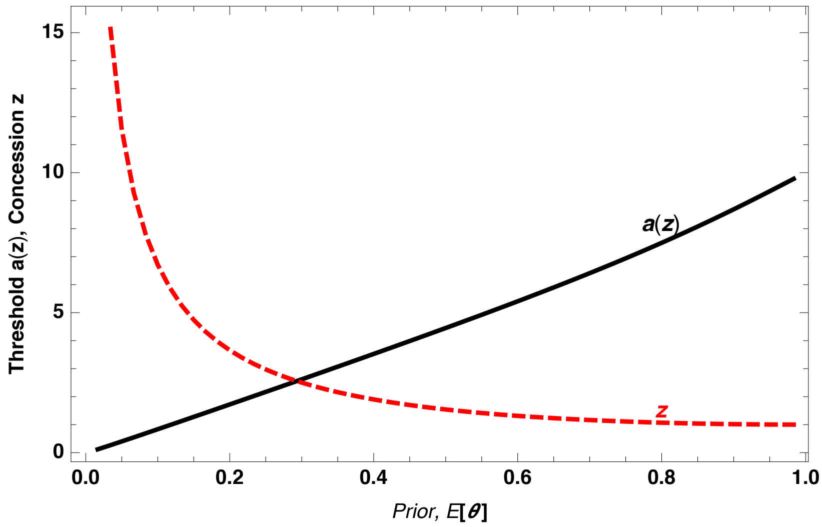

likely to occur, as on average, F's private signal is likely to be weaker when the prior $E[\theta]$ is small.

Quality of intelligence The basis of the quid pro quo is that F has some private information about the strength of L's regime. The number of trials, n, may be interpreted as the quality of F's intelligence. Both leaders' payoffs are increasing in n: the visit's deterrent value is increasing in the precision of the information that guided F's decision to conduct the visit, enhancing L's survival prospects; and in turn, L is willing to pay more for the visit, improving F's payoff. F's improved intelligence also means that F can better avoid visits with leaders that will be soon be removed, and avoid the costs associated with such a diplomatic misstep. An empirical implication is that visits with a US President—whose decisions are informed by $80 billion worth of annual intelligence gathering (DeVine 2019)—are far more valuable to recipient leaders than are visits with, say, a Canadian Prime Minister, who has no formal intelligence apparatus of his own (Robinson 2009). This relative valuation is not a function of the the countries' relative prestige or influence, but rather of the quality of private information that their leaders have access to. If we suppose that leaders face similar travel and opportunity costs for foreign visits (similar τ and ρ), then US Presidents should travel more than Canadian leaders because they can extract larger concessions in return.

7 Conclusions

We examine the bargaining that surrounds top-level diplomatic exchanges, and the impact of those exchanges on domestic political competition. A symbolic demonstration of support from one leader to another is valuable for the recipient because of the information it communicates to a potential challenger. This gives the incumbent an incentive to offer a concession to the foreign leader in exchange for a supportive diplomatic signal, but such a payment undermines the visit's signaling value. We show that mutually beneficial quid pro quo diplomacy can occur when negotiations occur behind closed doors. The price of a visit depends on the bargaining structure, the expected likelihood that a foreign leader wants a favor, the transparency of the concession, the recipient leader's prior regime strength, and the quality of the intelligence informing the foreign leader's decision. Provided there is a degree of opacity surrounding negotiations, symbolic public diplomacy can be informative even when it is purchased.

References

Alt, James E., Randall L. Calvert and Brian D. Humes. 1988. "Reputation and Hegemonic Stability: A Game-Theoretic Analysis." *American Political Science Review* 82(2):445–466.

Andersen, Thomas Barnebeck, Thomas Harr and Finn Tarp. 2006. "On US politics and IMF lending." *European Economic Review* 50(7):1843–1862.

Banks, Jeffrey. 1991. *Signaling games in political science.* Vol. 46 Psychology Press.

BBC News. 2007. "Gaddafi visit seals French deals.". http://news.bbc.co.uk/2/hi/africa/7135788.stm.

Bernstein, Richard. 2003. "Bush Visit Will Lift Poland to Status of Special Friend." The New York Times. https://www.nytimes.com/2003/05/29/world/bush-visit-will-lift-poland-to-status-of-special-friend.html.

DeVine, Michael E. 2019. "Intelligence Community Spending: Trends and Issues." *Congressional Research Service* 7-5700(R44381).

Fang, Songying. 2008. "The informational role of international institutions and domestic politics." *American Journal of Political Science* 52(2):304–321.

Gogia, Giorgi. 2020. "Political Graffiti Behind Bogus Jailing in Azerbaijan." Human Rights Watch. https://www.hrw.org/news/2020/02/18/political-graffiti-behind-bogus-jailing-azerbaijan.

Goldsmith, Benjamin, Yusaku Horiuchi and Kelly Matush. 2020. "Public Diplomacy Increases Foreign Public Approval: A Regression Discontinuity Analysis." *Working paper* . https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3566347.

Groseclose, Tim and Nolan McCarty. 2001. "The politics of blame: Bargaining before an audience." *American Journal of Political Science* pp. 100–119.

Holmes, Marcus and Keren Yarhi-Milo. 2016. "The psychological logic of peace summits: How empathy shapes outcomes of diplomatic negotiations." *International Studies Quarterly* 61(1):107–122.

Huseynov, Emin. 2016. "Freeze the Dictator Out." Foreign Policy. https://foreignpolicy.com/2016/03/30/freeze-the-dictator-out-azerbaijan-aliyev-obama-nuclear-security-summit/.

Lebovic, James H and Elizabeth N Saunders. 2016. "The diplomatic core: The determinants of high-level us diplomatic visits, 1946–2010." *International Studies Quarterly* 60(1):107–123.

Malis, Matt and Alastair Smith. 2019. "A global game of diplomacy." *Journal of Theoretical Politics* 31(4):480–506.

Malis, Matt and Alastair Smith. 2021. "State Visits and Leader Survival." *American Journal of Political Science* 65(1):241–256.

Matush, Kelly. 2020. "Harnessing Backlash: Why do leaders antagonize foreign states?" *Working paper* .

McAuley, James. 2018. "France's Sarkozy detained over allegations of taking money from Libya's Gaddafi.". The Washington Post.

McManus, Roseanne W. 2018. "Making it personal: The role of leader-specific signals in extended deterrence." *The Journal of Politics* 80(3):982–995.

Myrick, Rachel and Jeremy Weinstein. 2020. "Making Sense of Human Rights Diplomacy: Symbolism or Concrete Impact?" *Working paper* .

Nitsch, Volker. 2007. "State visits and international trade." *The World Economy* 30(12):1797–1816.

Office of the Vice President. 2016. "Readout of Vice President Biden's Meeting with President Ilham Aliyev of Azerbaijan.". https://obamawhitehouse.archives.gov/the-press-office/2016/03/31/readout-vice-president-bidens-meeting-president-ilham-aliyev-azerbaijan.

O'Grady, Siobhan. 2016. "Meet Ali Bongo Ondimba, Obama's Man in Africa." Foreign Policy. http://foreignpolicy.com/2016/04/05/meet-ali-bongo-ondimba-obamas-man-in-africa/.

Ostrander, Ian and Toby J Rider. 2019. "Presidents Abroad: The Politics of Personal Diplomacy." *Political Research Quarterly* 72(4):835–848.

Perlroth, Andres. 2019. "Posturing in Bargaining to Influence Outsiders." *Working paper* .

Permanent Select Committee on Intelligence. 2019. "Impeachment Inquiry: Ambassador William B. Taylor And Mr. George Kent." U.S. House of Representatives. https://judiciary.house.gov/uploadedfiles/2019-11-13_transcript_of_kent_taylor_hearing.pdf.

Robinson, Paul. 2009. "The Viability of a Canadian Foreign Intelligence Service." *International Journal* 64(3):703–716.
URL: *http://www.jstor.org/stable/40542197*

Savage, Charlie and Josh Williams. 2019. "Read the Text Messages Between U.S. and Ukrainian Officials." The New York Times. https://www.nytimes.com/interactive/2019/10/04/us/politics/ukraine-text-messages-volker.html.

Shadmehr, Mehdi and Raphael Boleslavsky. Forthcoming. "International Pressure, State Repression and the Spread of Protest." *Journal of Politics* .

Stasavage, David. 2004. "Open-door or closed-door? Transparency in domestic and international bargaining." *International organization* 58(4):667–703.

The New York Times. 2007. "Britain and Libya unveil energy and arms deals.". https://www.nytimes.com/2007/05/30/world/africa/30iht-30libya.5922646.html.

Vreeland, James Raymond and Axel Dreher. 2014. *The political economy of the United Nations Security Council: money and influence.* Cambridge University Press.

8 Appendix

8.1 Bayesian Updating

Here we more explicitly derive C's Bayesian updating upon observing the occurrence or absence of a visit, as summarized in Section 2.4 of the main text.

The prior belief on θ is the Beta distribution with parameters α and β, hence the pdf is $g(\theta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}$. Given the strategy that F visits if and only if $A \geq a$, we characterize C's beliefs following visit and non-visit. Given regime strength θ, the probability density with which $A = x$ is

$$p(x|\theta) = \frac{\Gamma(n+1)}{\Gamma(x+1)\Gamma(n-x+1)}\theta^x(1-\theta)^{(n-x)}$$

where we have stated the standard binomial coefficients $\begin{pmatrix} n \\ x \end{pmatrix} = \frac{n!}{x!(n-x)!}$ in terms of the gamma function since we do not restrict A to integers.

Averaging over all possible values of θ, let

$$\begin{aligned} p(x) &= \int_0^1 p(x|\theta)g(\theta)d\theta = \int_0^1 \frac{\Gamma(n+1)}{\Gamma(x+1)\Gamma(n-x+1)}\theta^x(1-\theta)^{(n-x)}\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}d\theta \\ &= \frac{\Gamma(n+1)\Gamma(x+\alpha)\Gamma(n-x+\beta)\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)\Gamma(x+1)\Gamma(n-x+1)\Gamma(n+\alpha+\beta)} \end{aligned}$$

and integrating over signals $A \leq a$ let

$$P(a) = \int_0^a p(x)dx = \int_0^a \frac{\Gamma(n+1)\Gamma(x+\alpha)\Gamma(n-x+\beta)\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)\Gamma(x+1)\Gamma(n-x+1)\Gamma(n+\alpha+\beta)}da \tag{10}$$

Therefore, $Pr(A < a) = \frac{P(a)}{P(n)}$ and since $E[\theta|A = x] = \frac{x+\alpha}{n+\alpha+\beta}$, we can express C's expectations of regime weakness, that is $(1-\theta)$, as

$$E[1-\theta|A < a] = \frac{\int_0^a \frac{\beta+n-x}{\alpha+\beta+n} p(x)dx}{\int_0^a p(x)dx} \tag{11}$$

Likewise, if C observes $A \geq a$,

$$E[1-\theta|A \geq a] = \frac{\int_a^n \frac{\beta+n-x}{\alpha+\beta+n} p(x)dx}{\int_a^n p(x)dx} \tag{12}$$

Differentiation of equations 11 and 12 yields

$$\frac{\partial E[1-\theta|A < a]}{\partial a} = \frac{p(a)}{P(a)} \left(E[1-\theta|A = a] - E[1-\theta|A < a]\right) < 0 \tag{13}$$

and

$$\frac{\partial E[1-\theta|A \geq a]}{\partial a} = \frac{p(a)}{P(n) - P(a)} \left(E[1-\theta|A \geq a] - \frac{n-a+\beta}{n+\alpha+\beta}\right) < 0 \tag{14}$$

8.2 Timing of Signals

The game form specified that F observed the signal A after concluding negotiations and before deciding whether to visit. An equally valid assumption might be that F learned the signal before negotiations, or that F learned partly before and partly afterwards. As it turns out, for the game with closed-door bargaining, equilibria under these alternative assumptions are behaviorally equivalent to the results reported in the main text. This equivalence is straightforward to see when L is proposer. Whether F accepts the proposal and visits only if $A \geq a(z)$, or F has prior knowledge of A and rejects agreements if $A < a(z)$, the path-of-play behavior is equivalent. Visits occur only when $A \geq a(z)$.

Likewise, when F is proposer, the timing of F information acquisition does not matter. In equilibrium, F demands the maximum that L will pay for a visit. F has no incentive to demand a smaller concession based upon a strong signal. The challenger does not observe the negotiations so it is observationally equivalent whether F demands the maximum concession that F will accept and then visits if $A \geq a(z)$, or if F knows A in advance and only demands the maximum concession when $A \geq a(z)$. From C's perspective, a visit indicates that L paid the maximum concession and that $A \geq a(z)$, and a non-visit indicates either $S = 0$ or F wanted a concession but $A < a(z)$. Thus the results are robust to alternative assumptions regarding the timing of when F learns about L's strength.

8.3 Open vs. Closed Bargaining, and Intermediate Transparency

The analysis considers two distinct bargaining protocols. Under open bargaining, C observes salience S, concession offer z, and whether the deal is accepted or rejected. Under closed bargaining, C observes none of these factors, but instead only probabilistically observes z in the event that a visit occurs (with probability q, which we call "concession transparency"). Thus the open vs. closed distinction bundles together multiple features of the bargaining protocol.

However, we note that any cases of intermediate transparency are effectively nested within one of these two extreme cases. Suppose that C observed the concession offer but did not observe salience. Since no offers are made absent a salient opportunity, C trivially infers from observing the offer that $S = 1$. Conversely, suppose that C observed salience, but did not observe the offer. If $S = 0$, then C knows there is no offer to observe. If $S = 1$, then the results are characterized in Proposition 8.

8.4 Decision to Visit

The concessions z_0 and z_n characterize limits. The concession z_0 is sufficiently large as to induce F to visit for even the weakest possible signal; given such a large concession, the occurrence of a visit is uninformative. In contrast, the concession z_n is not large enough to induce F to visit even for the strongest possible signal; given such a small concession, the absence of a visit is uninformative.

Proof of Lemma 1: Given C's belief that visits imply $A \geq a$, we can write F's payoff from visiting (equation (2)) given signal $A = a$ as

$$V(a,z) \quad = \quad z - \tau - ((1-r)z + \rho)\ \frac{n-a+\beta}{\alpha+\beta+n}\ E[1-\theta|A \geq a] \tag{15}$$

In equilibrium, $V(a,z) = 0$. From total differentiation of $V(a,z) = 0$, for $a \in (0,n)$,

$$\frac{da}{dz} \quad = \quad \frac{1 - \frac{n-a+\beta}{n+\alpha+\beta}E[1-\theta|A \geq a](1-r)}{(z(1-r)+\rho)\left(E[1-\theta|A \geq a]\frac{d\frac{n-a+\beta}{n+\alpha+\beta}}{da} + \frac{n-a+\beta}{n+\alpha+\beta}\frac{dE[1-\theta|A\geq a]}{da}\right)} < 0 \tag{16}$$

■

Proof of Propositions 1 and 2:

For $z \in (z_n, z_0)$, the equilibrium characterization follows directly from the implicit solution to $V(a,z) = 0$. F wants to visit if $A \geq a(z)$. Given this decision, C's beliefs are defined by Bayes' rule and $k_{nv}(z)$ and $k_v(z)$ are sequentially rational given these beliefs.

The characterization of the limiting cases requires tying down C's out-of-equilibrium beliefs. If F's strategy dictates that he always visits (i.e. visit if $A \geq 0$), then C's beliefs ($E[1-\theta|\text{non-visit}]$) are undefined by Bayes' rule in the contingency that F does not visit. Assumption 1 ensures that C infers non-visits to be associated with the worst possible signal, $E[1-\theta|\text{non-visit}] = E[1-\theta|A \leq 0] = \frac{\beta+n}{\alpha+\beta+n}$ and $k_{nv} = \frac{\beta+n}{\alpha+\beta+n}$. Likewise Assumption 1 ties down beliefs following a visit if F is never expected to visit ($a \geq n$) such that z_n is the limiting solution to $V(a,z) = 0$ as $a \to n$. The equilibrium is unique since the random variables A and K have no mass at any point. ■

The proof of Corollary 1 follows directly from total differentiation of Equation (2).

Proof of Corollary 2: Given the uniform prior ($\alpha = \beta = 1$), all signals are equally likely: $p(x) = \frac{1}{n}$ for all $x \in (0,n)$. Hence $E[1-\theta|A \geq a] = \int_a^n \frac{p(x)}{P(n)-P(a)}\frac{n+1-x}{n+2}dx = \frac{\frac{n+1-a}{n+2}+\frac{1}{n+2}}{2} = \frac{2-a+n}{4+2n}$. The result is an algebraic rearrangement of Equation (2): $z - \tau + ((1-r)z+\rho)\frac{2-a+n}{4+2n}\frac{n+1-a}{2+n} = 0$ ■

Proof of Proposition 3: Given concession $z \in (z_n, z_0)$, $Pr(visit|z) = 1 - \frac{P(a(z))}{P(n)}$ which is increasing in z: $\frac{dPr(visit|z)}{dz} = -\frac{p(a)}{P(n)}\frac{da(z)}{dz} > 0$. The second and third results follow from Equations (14) and (13).

For fixed z, $\frac{da(z)}{d\tau} = \frac{\partial V(a,z)}{\partial \tau}/\frac{\partial V(a,z)}{\partial a} > 0$, $\frac{da(z)}{d\rho} = \frac{\partial V(a,z)}{\partial \rho}/\frac{\partial V(a,z)}{\partial a} > 0$ and $\frac{da(z)}{dr} = \frac{\partial V(a,z)}{\partial r}/\frac{\partial V(a,z)}{\partial a} < 0$. ■

8.5 Open/Transparent Bargaining

Proof of Proposition 4: To simplify notation let $I_n = \int_0^a \frac{\beta+n-x}{\alpha+\beta+n} p(x)dx < P(a)$ and $I_v = \int_a^n \frac{\beta+n-x}{\alpha+\beta+n} p(x)dx < P(n) - P(a)$ Therefore, from Equation (5),

$$Pr(\text{Regime Change}|a(z)) = \frac{P(a)}{P(n)} \frac{I_n^2}{P(a)^2} + \frac{P(n)-P(a)}{P(n)} \frac{I_v^2}{(P(n)-P(a))^2}$$

and

$$E[1-\theta]^2 = \frac{(I_n + I_v))^2}{P(n)^2}$$

$$Pr(\text{Regime Change}|a(z)) - E[1-\theta]^2 = \frac{(I_n(P(a)-P(n)) + I_v P(a))^2}{P(a)P(n)^2(P(n)-P(a))} > 0$$

The inequality is strict because $I_n(P(a)-P(n)) + I_v P(a) < I_v(-I_n + P(a)) < 0$ so both numerator and denominator are positive. ■

Proposition 4 directly implies Proposition 5. L never agrees to or proposes any concession that results in a positive probability of a visit, since doing so increases her deposition risk.

9 Bargaining behind closed doors

Moving forward, it is useful to introduce some additional notation. We write $\chi = E[1-\theta] = \frac{\beta}{\alpha+\beta}$, $B(a) = E[1-\theta|A \geq a]$, $G(a) = E[1-\theta|A = a]$ and $H(a) = E[1-\theta|A < a]$.

Consider any $q \in [0,1]$, $\sigma \in (0,1)$ and suppose that the challenger conjectures that, if $S = 1$, then the negotiated concession is w. Supposing that the negotiated concession is actually z, we have that F will visit with probability $\frac{P(n)-P(a(z))}{P(n)}$. Should a visit occur, then with probability q, C will observe the concession z and consequently L's expected probability of deposition is $B(a(z))^2$; and with complementary probability $1-q$, the concession is not observed by C, who will infer $E[1-\theta|visit] = B(a(w))$ and L will infer $E[1-\theta|visit] = B(a(z))$ (since L knows the actual concession). Hence following a visit the expected probability of deposition is $qB(a(z))^2 + (1-q)B(a(w))B(a(z))$. If $S = 1$ and the negotiated concession is z, then with probability $\frac{P(a(z))}{P(n)}$, F does not visit. In this circumstance, L infers that $E[1-\theta|\text{non-visit}] = H(a(z))$. The challenger infers that the probability of a successful challenge is $E_c[1-\theta|\text{non-visit}, w]$, where, via Bayes' rule,

$$\begin{aligned} E_C[1-\theta|\text{non-visit}, w] &= Pr(S=1|\text{non-visit})H(a(w)) + (1-Pr(S=1|\text{non-visit}))\chi \\ &= \frac{P(a(w))\sigma H(a(w)) + (1-\sigma)P(n)\chi}{P(a(w))\sigma + (1-\sigma)P(n)} \end{aligned}$$

Given $S=1$, L's expected payoff from the agreement z is

$$\begin{aligned} U_L(z,w) &= -\overbrace{\frac{P(n)-P(a(z))}{P(n)}}^{\text{direct cost}} z + \Psi \left\{ 1 - \overbrace{\frac{(P(n)-P(a(z)))}{P(n)}}^{Pr(visit)} \overbrace{(qB(a(z))^2 + (1-q)B(a(w))B(a(z)))}^{\text{deposition risk}|\text{ visit}} \right. \\ &\left. - \underbrace{\frac{P(a(z))}{P(n)}}_{Pr(\text{non-visit})} \underbrace{H(a(z))\, E_C[1-\theta|\text{non-visit}, w]}_{\text{deposition risk}|\text{non-visit}} \right\} \end{aligned} \tag{17}$$

If no agreement is reached (or equivalently, $z < z_n$), then L's payoff is

$$U_L(z < z_n, w) = \Psi\,(1 - \chi\, E_C[1-\theta|\text{non-visit}, w]) \tag{18}$$

Proof of Proposition 6: Comparing (17) and (18), L will accept z if

$$\frac{z}{\Psi} \leq \frac{(\chi P(n) - H(a(z))P(a(z)))}{P(n)-P(a(z))} E_C[1-\theta|\text{non-visit}, w] - B(a(z))^2$$

F demands the largest concession that L will accept, which at the equilibrium condition $w=z$ implies (6). That the concession is increasing in σ follows from $E_C[1-\theta|\text{non-visit}, z]$ being increasing in σ, which implies that RHS of (6) is increasing in σ. As $\Psi \to \infty$, the LHS of (6) goes to 0 and as $z \to z_0$, $Q \to \chi$, $B(a(z)) \to \chi$ and $P(a(z)) \to 0$, so the RHS of (6) $\to 0$. ■

Proposition 6 holds as provided there is some possibility that F does not want a favor ($\sigma < 1$). However, when it is certain that F wants a favor $\sigma = 1$, then F can extract more than z_0 concessions, as we characterize below in proposition 8. The difference between the cases arises because when $\sigma = 1$ as $a \to 0$, $E_C[1-\theta|\text{non-visit}, w=z] \to H(0) > \chi$. However, if C believes that F wants a favor with probability $\sigma < 1$, then as $a \to 0$, $E_C[1-\theta|\text{non-visit}, w=z] \to \chi$.

Proposition 8 *If the Challenger is certain that a salient opportunity exists (i.e.*

$\sigma = 1$), then on the equilibrium path F demands the largest concession such that equation 19 holds and L accepts such an offer.

$$\frac{z}{\Psi} \leq \frac{(\chi P(n) - H(a(z))P(a(z)))}{P(n) - P(a(z))} H(a(z)) - B(a(z))^2 \tag{19}$$

If office holding is the dominant motive, $\Psi \geq \frac{(\alpha+\beta)^2(\alpha+\beta+n)(\tau(\alpha+\beta)(\alpha+\beta+n)+\beta\rho(\beta+n))}{\alpha\beta n(\alpha^2+2\alpha\beta+n(\alpha+\beta r)+\beta^2 r)}$, then F demands $z = \Psi\left(\frac{\alpha\beta n}{(\alpha+\beta)^2(\alpha+\beta+n)}\right)$ (which is the largest possible demand such that equation 19 holds given $a(z) = 0$). L accepts such an offer and given this deal F visits for any $A \geq 0$.

Proof of Proposition 8: When $\sigma = 1$, $E_C[1 - \theta|\text{non-visit}, w] = H(a(w))$. Evaluated at $w = z$, (6) implies (19). As $a(z) \to 0$, $B(a(z)) \to \frac{\beta}{\alpha+\beta}$, $H(a(z)) \to \frac{n+\beta}{n+\alpha+\beta}$ and the RHS of (19) is $\Psi\left(\frac{\alpha\beta n}{(\alpha+\beta)^2(\alpha+\beta+n)}\right)$. The condition on Ψ is sufficient to ensure that (19) holds for $z = z_0$. ■

9.1 L as proposer

Proof of Proposition 7: If L and F agree to concession z and C conjectures that the concession is w when he does not observe it, then F's payoff is

$$\begin{aligned} U_L(z,w) &= -\frac{(P(n) - P(a(z)))}{P(n)} z + \Psi\left\{1 - \frac{(P(n) - P(a(z)))}{P(n)} \left(qB(a(z))^2 + (1-q)B(a(w))B(a(z))\right)\right. \\ &- \left.\frac{P(a(z))}{P(n)} \frac{H(a(z))(P(a(w))\sigma H(a(w)) + (1-\sigma)P(n)\chi)}{P(n)(P(a(w))\sigma + (1-\sigma)P(n))}\right\} \end{aligned} \tag{20}$$

The first term corresponds to the direct cost. The second term correspond to L's probability of surviving in office.

F can do no better than accept any offer and visit if and only if $A \geq a(z)$ so to find L's best offer we need to maximize (20) with respect to z. The derivative of (20) with respect to z, evaluated at the equilibrium condition, $w = z$, is

$$\begin{aligned} \frac{dU_L(z,w)}{dz}\Big|_{w=z} &= \overbrace{\frac{a'(z)P'(a(z))}{P(n)}\Psi\left[-\frac{G(a(z))(P(n)\chi(1-\sigma) + P(a(z))H(a(z))\sigma)}{P(n)(1-\sigma) + P(a(z))\sigma} + B(a(z))(G(a(z))(1+q) - B(a(z))q)\right]}^{\text{office holding effect}} \\ &+ \underbrace{z\frac{a'(z)P'(a(z))}{P(n)} - \frac{P(n) - P(a(z))}{P(n)}}_{\text{direct cost}} \end{aligned} \tag{21}$$

where we substituted for $H'(a(z))$ and $B'(a(z))$ using (13) and (14).

Focusing on office holding benefits, as $\Psi \to \infty$,

$$\frac{1}{\Psi}\frac{dU_L(z,w)}{dz}|_{w=z} = \frac{a'(z)P'(a(z))}{P(n)}\left[-\frac{G(a(z))(P(n)\chi(1-\sigma)+P(a(z))H(a(z))\sigma)}{P(n)(1-\sigma)+P(a(z))\sigma} + B(a(z))(G(a(z))(1+q)-B(a(z))q)\right]$$

The first order condition implies (7) in Proposition 7. Next we show that this FOC is a maximum. The second order condition is proportional to the derivative of LHS of (7). In particular,

$$\frac{1}{\Psi}\frac{d^2U_L(z,w)}{dz^2}|_{w=z} = \frac{a'(z)^2P'(a(z))}{P(n)G(a(z))(P(n)-P(a(z)))} \times \left(qB(a(z))^2G'(a(z))(P(n)-P(a(z))) - 2qG(a(z))(B(a(z))-G(a(z)))^2P'(a(z))\right) < 0$$

This condition ensures that the FOC characterizes L's best offer. ■

Corollary 3 follows direct from total differentiation of (7).

Proof of Corollary 4: For the uniform distribution the first order condition (equation 7) is

$$\frac{a(z)^2(q+2) - 2a(z)(n+1)(q+1) + n(n+2)q}{4(n+2)^2} = 0$$

and equation 9 is the solution to this equation. To obtain the results on concession size, substitute $a(z)$ into equation 2.

The derivative $\frac{d\frac{a(z)}{n}}{dn} = \frac{n-(q+1)\left(\sqrt{n(n+2)+(q+1)^2}-q-1\right)}{n^2(q+2)\sqrt{n(n+2)+(q+1)^2}} < 0$ is negative. ■