Appendix

We present various additional results and extensions in the appendix. In Appendix A.1, we fully characterize the equilibrium, including a proof for Proposition 1 and a formal description of the strategies when conflict occurs along the equilibrium path. In Appendix A.2, we impose a lower bound on the per-period offer. In Appendix A.3, we model path-dependent shocks. In Appendix A.4, we model shocks in the cost of fighting (rather than the probability of winning). In Appendix A.5, we model endogenous institutional reform.

A.1 Full Equilibrium and Conflictual Paths of Play

We first prove Proposition 1. We then finish the equilibrium characterization of the baseline model by analyzing the parameters in which conflict occurs along the path of play, and discuss the per-period probability of conflict along conflictual paths of play. Note that a Markovian strategy for the ruler specifies a bargaining offer $x(p_t)$ as a function of the current-period threat p_t, and a Markovian strategy for the challenger specifies a response function as a function of p_t and x_t.

Proof of Proposition 1. The proof of existence and uniqueness proceeds in six steps. We first provide some properties of any Markovian equilibrium while treating the continuation value for the challenger as fixed (Steps 1–3), and then show that these properties uniquely pin down that the continuation value must be the expected value of fighting in the next period (Step 4). This enables us to establish the optimal per-period offer in a peaceful MPE (Step 5) and that the MPE is peaceful if and only if these offers are feasible for all draws of p_t (Step 6).

Step 1. In any MPE, the challenger accepts with probability 1 any offer that satisfies $x_t \geq \frac{p_t(1-\mu)}{1-\delta} - \delta V^C$ (see Equation 1) and accepts with probability 0 otherwise. For this and the next two steps, we denote as V^C the challenger's continuation value given the strategy profile (which is not a function of past play, given the Markovian restriction).

Proof: By construction, the challenger is strictly better off accepting if $x_t > \frac{p_t(1-\mu)}{1-\delta} - \delta V^C$ and rejecting if $x_t < \frac{p_t(1-\mu)}{1-\delta} - \delta V^C$. If the challenger were to not accept with probability 1 when the equality is met, then the ruler could profitably deviate to an infinitesimally larger offer, and therefore lacks a best response. So in any MPE the challenger must always accept $x_t = \frac{p_t(1-\mu)}{1-\delta} - \delta V^C$.

Step 2. Suppose $1 + \delta V^C \geq \frac{p_t(1-\mu)}{1-\delta}$. Then in any MPE, the ruler never offers an amount $x_t > \frac{p_t(1-\mu)}{1-\delta} - \delta V^C$.

Proof: The ruler's future continuation value is identical for any offer such that $x_t \geq \frac{p_t(1-\mu)}{1-\delta} - \delta V^C$ because, as the previous step established, the challenger will accept any such offer with probability 1. Consequently, any offer satisfying this inequality affects only the ruler's current-period payoff. The ruler can profitably deviate from any strategy in which the offer strictly satisfies the inequality because an infinitesimally smaller offer would increase the ruler's contemporaneous consumption without triggering the challenger to fight. Finally, the present assumption that $1 + \delta V^C \geq \frac{p_t(1-\mu)}{1-\delta}$ implies $x_t = \frac{p_t(1-\mu)}{1-\delta} - \delta V^C$ is a feasible offer.

Step 3. Suppose $1 + \delta V^C < \frac{p_t(1-\mu)}{1-\delta}$. Then in any MPE, the challenger rejects any feasible offer.

Proof: Follows directly from Step 1 and the assumed upper bound $x_t \leq 1$.

Step 4. In any MPE, $V^C = \frac{\bar{p}(1-\mu)}{1-\delta}$.

Proof: We demonstrate that $\frac{\bar{p}(1-\mu)}{1-\delta}$ comprises both a lower and upper bound for V^C, hence yielding the desired equality.

For the lower bound, it must be the case that $V^C \geq \frac{\bar{p}(1-\mu)}{1-\delta}$ because the challenger can always choose to fight in the next period. This yields in expectation the amount on the right-hand side of this inequality.

For the upper bound, along an arbitrary equilibrium path, we can assign a mixed probability with which the challenger accepts, $\alpha(p_t)$, for each state of the world (which also takes into account the ruler's optimal offer for that state). Any rejected offer yields $\frac{p_t(1-\mu)}{1-\delta}$. From Steps 2 and 3, we know that any offer that is accepted is bounded from above such that $x_t \leq \frac{p_t(1-\mu)}{1-\delta} - \delta V^C$. Thus, if we set accepted offers to this upper bound, we have an upper bound for the continuation value:

$$V^C \leq \int_{p^{\min}}^{p^{\max}} \Bigg[\overbrace{\alpha(p_t)}^{\text{Accept}} \Bigg(\underbrace{\frac{p_t(1-\mu)}{1-\delta} - \delta V^C}_{\text{Current period}} + \underbrace{\delta V^C}_{\text{Future periods}} \Bigg) + \overbrace{\big(1-\alpha(p_t)\big)}^{\text{Reject}} \frac{p_t(1-\mu)}{1-\delta} \Bigg] dF(p)$$

$$= \frac{\bar{p}(1-\mu)}{1-\delta}.$$

Thus, the upper and lower bounds for V^C are equal, which completes the proof.

Step 5. Suppose $\frac{(p_t - \delta\bar{p})(1-\mu)}{1-\delta} \leq 1$. Then in any MPE, the ruler proposes $x_t = x^*(p_t) \equiv \frac{(p_t-\delta\bar{p})(1-\mu)}{1-\delta}$.

Proof: Substituting V^C from Step 4 into the inequality from Step 1 demonstrates that the challenger will accept $x_t = \frac{(p_t-\delta\bar{p})(1-\mu)}{1-\delta}$, which is the term from Equation (4). Furthermore, because of the present assumption $\frac{(p_t-\delta\bar{p})(1-\mu)}{1-\delta} \leq 1$, this offer is feasible. To show this is the unique offer in any MPE when $\frac{(p_t-\delta\bar{p})(1-\mu)}{1-\delta} \leq 1$ is satisfied, we show that offering any other amount makes the ruler strictly worse off.

To see that a higher offer makes the ruler strictly worse off, in Step 2 we established that the ruler never benefits from offering an amount $x_t > \frac{p_t(1-\mu)}{1-\delta} - \delta V^C$. Substituting in the value of V^C solved for in Step 4 yields the desired inequality.

To see that a lower offer makes the ruler strictly worse off, we begin by observing that the challenger will reject any lower offer (see Step 1 while substituting in V^C from Step 4). Therefore, a deviation by the ruler will yield a lifetime expected utility for herself of $\frac{(1-p_t)(1-\mu)}{1-\delta}$. To demonstrate that this deviation is strictly unprofitable, we need to demonstrate that the ruler can ensure herself a higher payoff from offering $x_t = \frac{(p_t-\delta\bar{p})(1-\mu)}{1-\delta}$ and thereby securing acceptance in the present period. Following a period of acceptance, we can bound the ruler's future continuation value from below at $\frac{(1-\bar{p})(1-\mu)}{1-\delta}$ because the ruler can always trigger the challenger to fight in the next period. Thus the claim requires showing that the ruler's lower-bound lifetime expected utility from securing agreement in the current period strictly exceeds her utility to fighting now:

$$\underbrace{1-\frac{(p_t-\delta\bar{p})(1-\mu)}{1-\delta}}_{\text{Peace now}}+\underbrace{\delta\frac{(1-\bar{p})(1-\mu)}{1-\delta}}_{\text{Conflict in next period}}>\underbrace{\frac{(1-p_t)(1-\mu)}{1-\delta}}_{\text{Conflict now}}.$$

This reduces to $\delta(1-\mu)>0$, which always holds.

Step 6. There is a unique peaceful MPE if and only if $\frac{(p^{\max}-\delta\bar{p})(1-\mu)}{1-\delta}\leq 1$.

Proof: To prove "if," note that $x^*(p_t)$ strictly increases in p_t. Therefore, if $x^*(p^{\max})\leq 1$, then this inequality holds for all p_t. Per the proceeding steps, the optimal-offer function for the ruler and the optimal-response function for the challenger are each unique, and the challenger accepts in every period.

To prove "only if," if the inequality is violated, then the challenger rejects any feasible offer when $p_t=p^{\max}$. This follows from Step 3 after substituting in the term for V^C established in Step 4. Consequently, conflict occurs along the equilibrium path. □

We now provide a full characterization of equilibrium strategies when the condition in Proposition 1 is violated.

Proposition A.1 (Conflictual equilibrium). *If* $\frac{(p^{max}-\delta\bar{p})(1-\mu)}{1-\delta}>1$*, then there is a unique class of payoff-equivalent MPE in which conflict occurs along the path of play. In these MPE, there is a unique* $p^*\in(p^{min},p^{max})$ *such that (i) when* $p_t\leq p^*$*, the strategies are the same as in Proposition 1, and (ii) when* $p_t>p^*$*, then the challenger rejects all offers and the ruler's strategy can involve making any offer.*

Proof The proof of Proposition 1 provides most of the elements needed to establish existence and uniqueness. For the present proof, we formally define p^* as:

$$\frac{\big(p^*-\delta\bar{p}\big)(1-\mu)}{1-\delta}=1\implies p^*=\frac{1-\delta}{1-\mu}+\delta\bar{p}.$$

For the bounds, the upper bound $p^* < p^{\max}$ follows from the present assumption that $\frac{(p^{\max}-\delta\bar{p})(1-\mu)}{1-\delta} > 1$. The lower bound $p^* > p^{\min}$ requires $\frac{(p^{\min}-\delta\bar{p})(1-\mu)}{1-\delta} < 1$. Algebraic rearrangement yields $p^{\min}(1-\mu) < \bar{p}(1-\mu) + (1-\delta)(1-\bar{p}(1-\mu))$, which is true because $p^{\min} < \bar{p}$.

The unique optimality of the challenger's accept/fight decisions follows immediately from the construction of p^* and from the steps in the proof of Proposition 1; as do the offers from the ruler when $p \leq p^*$. When $p > p^*$, all offers are rejected, and therefore the ruler is indifferent among all offers. □

How challenger strength affects the per-period probability of conflict Throughout the analysis in the article, when we assess prospects for conflict, we mean prospects for an equilibrium in which conflict occurs along the path of play. Here we extend the analysis by considering how challenger strength affects outcomes within the set of parameter values in which conflict occurs along the equilibrium path. Along a conflictual equilibrium path, the per-period probability of conflict (assuming none has occurred previously) is the probability of drawing $p_t > p^*$. Since the cumulative distribution function of p_t is F, we can write this $1 - F\left(\frac{1-\delta}{1-\mu} + \delta\bar{p}\right)$.

Increasing challenger strength changes two terms in this expression: $\bar{p}$ and the F function. Suppose we define an increase in challenger strength as a uniform upward shift in the probability of winning a conflict, such that this probability is $p_t + d$ for some constant $d > 0$. In this formulation, p_t is the "baseline" probability of winning, which still follows distribution F, and d is the increase in this baseline threat. Thus, we can use the expressions from above while replacing p_t with $p_t + d$, and $\bar{p}$ with $\bar{p} + d$. Consequently, the per-period probability of conflict is $Pr\Big(p_t + d > \frac{1-\delta}{1-\mu} + \delta(\bar{p}+d)\Big) = 1 - F\left(\frac{1-\delta}{1-\mu} + \delta\bar{p} - (1-\delta)d\right)$. This term strictly increases in d. Therefore, conditional on conflict occurring along the equilibrium path, a stronger challenger (defined by a uniform shift) decreases the expected number of periods until conflict occurs. A uniform upward shift in threats improves the challenger's continuation value from accepting (because it gains higher average offers in the future) and from fighting (because it wins with higher probability). The latter term dominates the former term because it is not discounted by a period, as discussed following the statement of Proposition 6.

The binary-threat case permits us to explore the effects of a shift in the distribution function itself. One notion of a stronger challenger is a higher frequency of maximum-threat periods, expressed by q. In the text, we demonstrated that higher q increases the range of parameter values in which the equilibrium is peaceful. However, conditional on the equilibrium path featuring conflict, higher q in fact *raises* the per-period probability of conflict. A high value of q guarantees peace; it is straightforward to verify that the condition in Proposition 1 always holds in the binary case if $q = 1$. However, the cause of the higher average threat is that maximum-threat periods arise more frequently—which means that conflict is expected to occur sooner if that event ever occurs along the equilibrium path. Overall, the effect of q on the per-period probability of conflict is non-monotonic: positive and strictly increasing until it drops to 0.

Figure A.1: Expected Time Until Conflict in Binary Threats Model

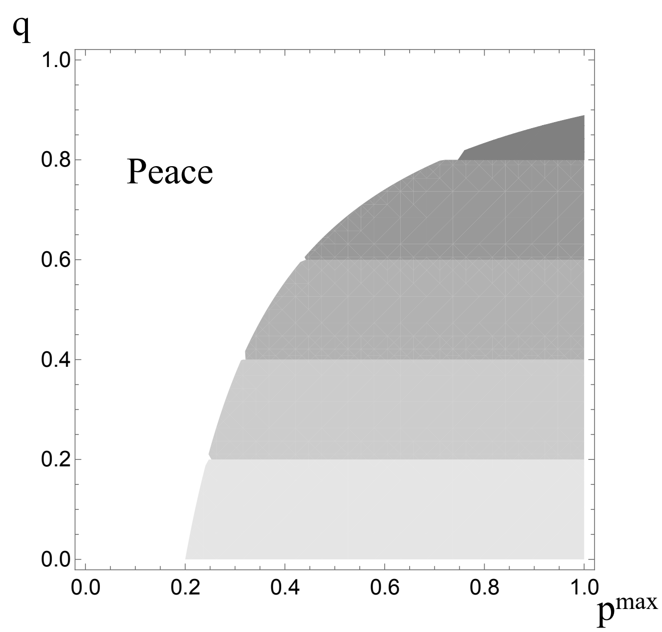

Parameter values: $\delta = 0.9$, $\mu = 0.5$, $p^{\min} = 0$.

We can see this visually in Figure A.1. It has the same parameter values and general setup as in Figure 1 except now we provide information on what happens in a conflictual path of play. The per-period probability of conflict is 0 in the white area (i.e., a peaceful path of play), and is positive in the gray areas (i.e., a conflictual path of play); and darker colors indicate a higher per-period probability of conflict. The non-monotonic effect of q is readily apparent: the total size of the conflict region is smaller for higher values of q, but conditional on conflict occurring along the equilibrium path, it is expected to occur sooner.

This finding highlights another twist in understanding the overall relationship between challenger strength and conflict. Depending on parameter values, a medium-sized increase in q can in fact make conflict *more imminent*, whereas a large increase in q eliminates conflict entirely.

A.2 Lower Bound on Offers

Here we extend the model to assume that the per-period offer must satisfy $x_t \in [\underline{x}, 1]$, for an exogenously specified $\underline{x} < 1$. A natural value to consider is $\underline{x} = 0$, that is, the ruler cannot demand net transfers away from the challenger, although the following results hold for more general values of $\underline{x}$. We derive these results under the specific case of binary challenger strength, while allowing strength to affect the minimum and maximum threats in addition to the probability of a maximum-threat period. Specifically, $p_t \in \{p^{\text{min}}, p^{\text{max}}\}$, with $q = Pr(p_t = p^{\text{max}})$. Let $x(p_t, \underline{x})$ be the offer made when the current-period threat is p_t and the lower bound on offers is $\underline{x}$. For the unbounded case we analyze in the text, we write $x(p_t, -\infty)$. At the end of this section, we comment on modeling a lower bound for the more general distribution of threats.

By Proposition 1, in any peaceful MPE, the offers in each period satisfy:

$$x^*(p^{\text{min}}, -\infty) = \frac{1}{1-\delta}\Big((1-\delta(1-q))p^{\text{min}} - \delta q p^{\text{max}}\Big)(1-\mu)$$
$$x^*(p^{\text{max}}, -\infty) = \frac{1}{1-\delta}\Big(p^{\text{max}}(1-\delta q) - \delta(1-q)p^{\text{min}}\Big)(1-\mu).$$

If $\underline{x} \leq x^*(p^{\text{min}}, -\infty)$, then the lower bound never binds and the analysis is equivalent to the unbounded case. At the other extreme, if $\underline{x} > p^{\max}(1-\mu)$, then the challenger accepts the basement offer even in a maximum-threat period.

If $\underline{x}$ is in-between these extremes, then along a peaceful equilibrium path, the ruler will offer $\underline{x}$ in a minimum-threat period and make a higher offer in a maximum-threat period. In such an equilibrium, the offer made in a maximum-threat period must make the challenger indifferent between accepting and not:

$$x^*(p^{\max}, \underline{x}) + \frac{\delta}{1-\delta}\Big(qx^*(p^{\max}, \underline{x}) + (1-q)\underline{x}\Big) = \frac{p^{\max}(1-\mu)}{1-\delta}$$

$$\implies x^*(p^{\max}, \underline{x}) = \frac{p^{\max}(1-\mu) - \delta\underline{x}(1-q)}{1-\delta(1-q)}.$$

Given the upper bound of 1 for an offer, a peaceful MPE requires $x^*(p^{\max}, \underline{x}) \leq 1$. The offer in a maximum-threat period decreases in $\underline{x}$ because higher basement spoils increase the challenger's average consumption in future periods. We can rearrange to show that $x^*(p^{\max}, \underline{x}) \leq 1$ if and only if:

$$\underline{x} \geq 1 - \frac{1 - p^{\max}(1-\mu)}{\delta(1-q)} \equiv \underline{x}^{\text{peace}}. \tag{A.1}$$

This threshold is strictly less than 1, which means it is always possible to set $\underline{x}$ high enough to induce a peaceful equilibrium path of play.

Finally, we point out that there is no reason to believe that the core insights would not extend for the more general distribution of threats in our baseline model. However, the general case is difficult to characterize analytically. Intuitively, whenever p_t is lower than some bound $\underline{p}$, the ruler will offer exactly $x_t = \underline{x}$, and for all other periods the ruler will offer a higher value of x_t that makes the challenger indifferent between accepting and fighting. This breaks the linear structure of the offers in the baseline case. The specific complication is that the threshold $\underline{p}$ is endogenous to anticipated outcomes along the future path of play. This makes it difficult to characterize clean comparative statics on key parameters such as challenger strength.

A.3 Path-Dependent States

Despite the generality of our baseline model, one stark assumption is that Nature draws threat levels independently across periods. A simple way to introduce path-dependent states is to assume in each period that the challenger's threat level is either identical to the level from the previous period or a new draw occurs from the main distribution $F(p; s)$. The probability that the threat level persists from the previous period is $r \in (0, 1)$. The main findings here are (1) more persistent threats make conflict less likely; and (2) if threats are sufficiently persistent, then stronger challengers are unambiguously harder to buy off.

In this extension, the continuation value depends on the current value of p_t. Let $V^C(p_t)$ be the continuation value for entering the next period when the current threat is p_t. We can write the indifference condition as:

$$x_t(p_t) = \frac{p_t(1-\mu)}{1-\delta} - \delta\Big(rV^C(p_t) + (1-r)V_n^C\Big), \tag{A.2}$$

where $V_n^C = \mathbb{E}[V^C(p_t)]$ is the continuation value if the threat is "new." We can write the continuation value with threat p_t as:

$$V^C(p_t) = x_t(p_t) + \delta\Big(rV^C(p_t) + (1-r)V_n^C\Big)$$

$$\implies V^C(p_t) = \frac{x_t(p_t) + \delta(1-r)V_n^C}{1-\delta r}.$$

Substituting this term back into Equation (A.2) yields:

$$x_t(p_t) = \frac{1}{1-\delta}p_t(1-\mu) - \delta\Big(r\frac{x_t(p_t) + \delta(1-r)V_n^C}{1-\delta r} + (1-r)V_n^C\Big)$$

$$\implies x_t(p_t) = \frac{1-\delta r}{1-\delta} p_t(1-\mu) - \delta(1-r)V_n^C. \tag{A.3}$$

Importantly, and as in our baseline analysis, this expression is linear in p_t. As a result, we can solve for V_n^C as follows:

$$\begin{aligned} V_n^C &= \mathbb{E}[x_t(p_t)] + \delta V_n^C \\ &= \frac{1-\delta r}{1-\delta} \bar{p}(1-\mu) - \delta(1-r)V_n^C + \delta V_n^C. \end{aligned}$$

Solving for V_n^C gives:

$$V_n^C = \frac{\bar{p}(1-\mu)}{1-\delta}. \tag{A.4}$$

Note that this expression is the same as in the baseline case without path dependence, $r = 0$. Substituting Equation (A.4) back into Equation (A.3) provides an explicit characterization of the offer in each period:

$$\begin{aligned} x_t(p_t) &= \frac{1-\delta r}{1-\delta} p_t(1-\mu) - \delta(1-r)\frac{1}{1-\delta}\bar{p}(1-\mu) \\ \implies x_t(p_t) &= \frac{1}{1-\delta}\Big((1-\delta r)p_t - \delta(1-r)\bar{p}\Big)(1-\mu). \end{aligned}$$

As $r \to 0$, we recover our baseline setup without path dependence. As $r \to 1$, threats do not change over time, and hence the optimal offer converges to the offer from a static version of the model, $(1-\mu)p_t$. This term is strictly less than 1, which means that any equilibrium path of play is peaceful. This is expected; the reason that fighting can occur along the equilibrium path in bargaining models with limited commitment is that threat levels fluctuate over time.

In general, peace is possible when:

$$\frac{1-\delta}{1-\mu} \geq \underbrace{p^{\max} - \delta\bar{p}}_{\tau(s)} - \delta r(p^{\max} - \bar{p}) \equiv \tau(s, r). \tag{A.5}$$

The first result is that more persistent threats make conflict less likely. As we can see in Equation (A.5), higher r makes the inequality true for a wider range of parameter values; and at $r = 0$ it collapses to Equation (5).

The second result is that if threats are sufficiently persistent, then stronger challengers are unambiguously harder to buy off. The inequality in Equation (A.5) is harder to sustain for a stronger challenger if $\tau(s, r)$ increases in s:

$$(1-\delta r)\frac{\partial p^{\max}}{\partial s} - \delta(1-r)\frac{\partial \bar{p}}{\partial s} > 0.$$

To yield the result, as $r \to 1$, the second term in the preceding expression approaches zero, whereas the first term approaches $(1-\delta)\frac{\partial p^{\max}}{\partial s}$. Consequently, $\frac{\partial p^{\max}}{\partial s} > 0$ implies that the preceding inequality must hold.

A.4 Fluctuating Costs of Conflict

In this section, we analyze a variant of the model in which the probability of winning is fixed but the cost of fighting fluctuates across periods. This more closely resembles the setup in Acemoglu and Robinson (2006), and the insights are qualitatively identical to our baseline model.

Suppose the probability of challenger victory is fixed at $p \in (0, 1]$ and the fraction of spoils that would permanently be destroyed by conflict is given by μ_t. We rule out the trivial case $p = 0$, in which it is immediately apparent that the ruler survives while offering nothing in each period. In each period, μ_t is drawn iid from a distribution $G(\mu)$ with minimum value $\mu^{\min}$, maximum value $\mu^{\max}$, and average value $\bar{\mu}$.

By an identical logic as in our baseline model, the optimal transfer in every period must satisfy:

$$x^*(\mu_t) = \frac{p(1-\mu_t)}{1-\delta} - \delta V^C.$$

In a peaceful MPE, the continuation value is written as follows. The first line is identical to the baseline setup except the integrand differs, and the final expression for V^C is identical except the average is taken over μ rather than p.

$$V^C = \frac{1}{1-\delta}\int_{\mu^{\min}}^{\mu^{\max}} \left(\frac{p(1-\mu)}{1-\delta} - \delta V^C\right) dG(\mu)$$

$$\implies V^C = \frac{p(1-\bar{\mu})}{1-\delta}.$$

Consequently, the optimal offer is:

$$x^*(\mu_t) = \frac{p}{1-\delta}\Big(1-\delta-(\mu_t-\delta\bar{\mu})\Big).$$

The condition for a peaceful MPE is that it is possible to buy off the challenger when conflict destroys the smallest share of the pie, or:

$$\frac{p}{1-\delta}\Big(1-\delta-(\mu^{\min}-\delta\bar{\mu})\Big) \leq 1.$$

This yields qualitatively identical comparative statics as the main analysis. If increasing challenger strength decreases the average amount destroyed by conflict but not the minimum amount, then this inequality is easier to meet, and so stronger challengers are easier to buy off peacefully. By contrast, if making the challenger stronger decreases $\mu^{\min}$ and $\bar{\mu}$ at an equal rate, then the opposite holds.

A.5 Endogenous Institutional Reform

In Appendix A.2, we extended the binary threat version of the model to incorporate an exogenous lower bound $\underline{x}$ on the ruler's per-period offer. Now we endogenize the choice of $\underline{x}$, which we interpret as strategic institutional reform. In each period, after Nature realizes the challenger's threat, the ruler chooses $\underline{x}_t \in [\underline{x}_{t-1}, 1]$, with the initial level corresponding to that in the baseline game, $\underline{x}_0 = -\infty$. This means that the institutional choice in any period is a dynamic state variable and creates a floor for the offer in all future periods; the ruler can subsequently choose to raise this floor, but not lower it. This choice could capture a wide range of institutional reforms, such as a power-sharing agreement, expanding the franchise, or civil rights protections.

We begin by presenting three preliminary results. First, if the inequality in Proposition 1 is met, then the ruler will not set $\underline{x}_t > -\infty$. A different choice would either have no impact on the outcome the game or would redistribute more surplus than needed to buy off the challenger. As a result, we focus on the case when the inequality in Proposition 1 is not met, and hence conflict will occur along the equilibrium path absent reform.

Second, the ruler never has a strict preference to reform institutions in a minimum-threat period. Doing so would deliver (weakly) more transfers to the challenger in a period in which it can already be induced to accept (see the proof for Proposition A.1), but has no impact on the ruler's ability to buy off the challenger in a maximum-threat period. In such a period, the ruler can instantaneously increase the basement level of transfers.

Third, if the ruler makes institutional reforms, they will be "large." Recall from Equation (A.1) that $\underline{x}^{\text{peace}}$ is the level of $\underline{x}_t$ at which the challenger is indifferent between accepting an offer of 1 and fighting in a maximum-threat period. This is the lowest level of $\underline{x}_t$ that induces a peaceful path of play. It is straightforward to rule out any finite choice $\underline{x}_t < \underline{x}^{\text{peace}}$ as the optimal level of institutional reform. Such a choice does not change the challenger's preference to fight in maximum-threat periods and simply delivers weakly more spoils to the challenger in minimum-threat periods, when it would accept anyway.

Given these preliminary results, we ask: in a maximum-threat period, if conflict would otherwise occur, will the ruler make institutional reforms sufficiently large to buy off the challenger? The following proves that the answer is always yes. We already know the ruler's lifetime expected utility if a conflict occurs in a maximum-threat period:

$$\frac{(1-p^{\text{max}})(1-\mu)}{1-\delta}. \tag{A.6}$$

Alternatively, upon choosing $\underline{x}_t \geq \underline{x}^{\text{peace}}$ but not subsequently choosing $\underline{x}_z > \underline{x}_t$ in any period $z > t$, the ruler's lifetime expected utility is:

$$1 - x^*(\underline{x}_t) + \frac{\delta}{1-\delta}\Big(q(1 - x^*(\underline{x}_t)) + (1-q)(1-\underline{x}_t)\Big), \tag{A.7}$$

where $x^*(\underline{x}_t)$ is the offer that makes the challenger indifferent between accepting and fighting in a maximum-threat period given institutions $\underline{x}_t$. This offer must satisfy:

$$x^*(\underline{x}_t) + \frac{\delta}{1-\delta}\Big(qx^*(\underline{x}_t) + (1-q)\underline{x}_t\Big) = \frac{p^{\text{max}}(1-\mu)}{1-\delta}. \tag{A.8}$$

Solving Equation (A.8) for $x^*(\underline{x}_t)$ and substituting back into Equation (A.7) yields a lifetime expected utility for the ruler of:

$$\frac{1-p^{\text{max}}(1-\mu)}{1-\delta}. \tag{A.9}$$

Finally, we compare Equations (A.6) and (A.9):

$$\frac{1-p^{\text{max}}(1-\mu)}{1-\delta} > \frac{(1-p^{\text{max}})(1-\mu)}{1-\delta}$$

which holds for any $\mu > 0$.

One notable attribute about the preceding proof is that conditional on making a large-enough institutional reform to induce peace, the ruler is in fact indifferent about the exact amount of institutional reform. There are a continuum of equilibrium choices in which the ruler chooses between a bit more institutional reform (yielding less consumption for herself in a minimum threat period) and offering somewhat fewer temporary transfers in a maximum threat period, and vice versa. We focus on the MPE with the minimum-necessary institutional reforms, which is consistent with microfoundations for such a choice posited in Castañeda Dower et al. (2018) and Powell (2021).

Along the equilibrium path, the ruler does not choose institutional reform until the first maximum-threat period, when she implements reform. Formally, the ruler optimally sets $\underline{x}_t = \max\{-\infty, \underline{x}_{t-1}\}$ in every minimum-threat period and $\underline{x}_t = \max\{\underline{x}^{\text{peace}}, \underline{x}_{t-1}\}$ in every maximum-threat period; the max function accounts for the inability to lower basement spoils below those chosen in previous periods.

Given this result, the comparative statics on s are identical to those in the baseline game. We simply replace the conflict region in Figure 1 with a "reform" region. In other words, the parameter values in the baseline model for which conflict would ensue is identical to the parameter values in the present extension for which institutional reform will occur.

Therefore, higher p^{max} increases the range of parameter values in which any institutional reform occurs. An additional result is that higher p^{max} also increases the extent of institutional reforms (conditional on any occurring). To establish this result, we differentiate $\underline{x}^{\text{peace}}$ (see Equation (A.1)) with respect to p^{max}. Increasing the challenger's opportunity cost to not fighting in a maximum-threat period bolsters the credibility of its demands for more institutional reform.