

The Guardianship Dilemma: Regime Security *through* and *from* the Armed Forces*

R. Blake McMahon

Branislav L. Slantchev

January 7, 2015

Forthcoming in the *American Political Science Review*

Abstract. Armed forces strong enough to protect the state also pose a threat to the state. We develop a model that distills this “Guardianship Dilemma” to its barest essentials, and show that the seemingly ironclad logic underlying our existing understanding of civil-military relations is flawed. Militaries contemplating disloyalty must worry about both successfully overthrowing the government *and* defeating the state’s opponent. This twin challenge induces loyalty as the state faces increasingly strong external threats, and can be managed effectively by rulers using a number of policy levers. Disloyalty can still occur when political and military elites hold divergent beliefs about the threat environment facing the state, since militaries will sometimes have less incentive to remain loyal than the ruler suspects. Consequently, it is not the need to respond to external threats that raises the risk of disloyalty — as conventional wisdom suggests — but rather uncertainty about the severity of these threats.

*Email: rmcmahon@ucsd.edu and slantchev@ucsd.edu. We thank Jesse Driscoll, Charles Han-
kla, Seth Hill, Allison Beth Hodgkins, Paul Johnson, and David Wiens for comments. Earlier ver-
sions were presented at 2014 meetings of the Midwest Political Science Association, American
Political Science Association, and ISSS-ISAC. Financial support from the National Science Founda-
tion is gratefully acknowledged by both McMahon (Graduate Research Fellowship) and Slantchev
(Grant SES-1153441). McMahon also thanks the Institute on Global Cooperation and Conflict for
its generous support.

Mercenary captains are either excellent men of arms or not: if they are, you cannot trust them because they always aspire to their own greatness, either by oppressing you, who are their patron, or by oppressing others contrary to your intention; but if the captain is not virtuous, he ruins you in the ordinary way.

Machiavelli, *The Prince*, XII

Rulers govern in an environment characterized by foreign and domestic threats, and must provide for their security if they are to survive in power. The state must therefore rely on a coercive force — one that specializes in dealing with foreign adversaries and another that focuses on internal ones or, as for most of history, one that could be used for either — an agent powerful enough to cope with these challenges, a “guardian” of the government. This existential imperative gives rise to one of the oldest dilemmas of governing, for a guardian strong enough to protect the government is also strong enough to impose its own preferences on the polity. Paradoxically, the attempt to cope with some types of threats can generate a new, and often very serious, threat. This *Guardianship Dilemma* can be resolved in two ways: the government either creates the forces it needs and takes its chances that they may turn on it or avoids that danger altogether by leaving itself exposed to the other threats. The more grave these other threats are, the more likely is the government to accept the risks of guardianship and opt for the creation of a military force (Huntington, 1957; Feaver, 1999; Svobik, 2012).

The unfortunate tradeoff for regime security implied by the Guardianship Dilemma is seen as a fundamental feature of civil-military relations, to the point where it has become an unstated assumption. As a result, most research on the subject explores the ways in which governments can maintain the necessary forces without running the risks of becoming their servant or getting eliminated altogether. The remedies suggested range from institutional checks and balances with countervailing forces to placing limits on budgets or the competence of military leaders, and from imposing organizational straight-jackets through manipulating the chain of command, recruitment, or inter-agency communications to the fostering of a separate apolitical professional culture in the military (Finer, 1988; Quinlivan, 1999; Pilster and Boehmelt, 2011; Egorov and Sonin, 2011).

But is the ancient logic underlying the Guardianship Dilemma compelling? We present a model of the interaction between the government and its military force that is starkly reduced to the barest essentials identified by the dilemma and show that, as commonly posed, the dilemma is flawed because it fails to account for the effect of the threat environment on the incentives of the guardians to interfere with politics. Because armed forces that intervene in politics must both successfully execute a coup *and* fend off other challengers, grave threats to the existing government can *induce* military loyalty. While the Guardianship Dilemma predicts that rulers are at greatest risk of a coup when some threat forces them to strengthen their mil-

itaries, we show that when these leaders are aware of the extent of this threat, it is, in fact, precisely the serious threat that permits them to pour resources into the military without fearing that it will become disloyal. This finding is consistent with the pathbreaking work of Desch (1999), who argues that large external threats help political leaders maintain control of the armed forces. In contrast to Desch, however, our model also reveals that when rulers know the threat's severity, it is possible to devise a combination of benefits that increase the military's value of remaining loyal and constraints that hinder its ability to fight, such that military loyalty is assured under all conditions. The dilemma is always resolved in favor of regime stability regardless of the size of the threat, although sometimes this happens at the expense of the military's effectiveness.

This indicates that to trigger the full power of the dilemma, another factor needs to be considered: something that prevents rulers from succeeding in the delicate balancing act between having an overweening military that might overthrow them and having an impotent one that might be useless against the threat. We argue that this necessary factor is asymmetric information about the threat environment, a type of uncertainty inherent in civil-military relations. When the military is better-informed about the extent of the threat facing the polity than the rulers — a discrepancy that could arise for a variety of reasons, as we explain below — the delicate balancing act can become well-nigh impossible.

We show that under general conditions rulers must end up with one of two unpleasant realities. If they decide that the threat is likely to be small, the military is endowed with just enough resources to deal with small threats. Whereas this ensures the military's loyalty in all circumstances, the military will be woefully unprepared if the politicians were too optimistic and the actual threat turns out to be large. If, on the other hand, rulers decide that the threat is likely to be large, they are hit with a double whammy: in their fear of a coup, rulers still handicap the military and reduce its effectiveness in dealing with the large threat, but because they are also fearful of the threat itself, rulers still endow the military with enough resources to induce its disloyalty if the threat happens to be small. In this case, the military is both ineffective against the serious threat and a danger to the regime if the rulers' estimates prove to be too pessimistic.

Thus, *the Guardianship Dilemma arises because of a mismatch between the military's strength and the threat it is supposed to deal with — if the military is underfunded, it will be loyal but deficient, and if it is over-funded, it will be effective but potentially disloyal — and the mismatch itself is caused by the divergent beliefs about the seriousness of the threat among the political and military leaders.* This divergence can be a product of the military's specialization in dealing with threats, which entails access to superior intelligence and information processing when it comes to estimating potential dangers to the polity. The closer the rulers get to the military's own estimates about the threat, the narrower the belief gap, and the weaker the dilemma.

Our model, reduced though it is, allows us to qualify some of the claims that are often made in studies of civil-military relations. For example, it is often argued that in making military appointments, non-democratic leaders prefer to select for loyalty rather than competence, and that as a result their armed forces are frequently quite ineffective on the battlefield. We explain why this line of reasoning is problematic. Competence and resources are both means to an end — higher probability of success in a military confrontation, be it with the threat or in a coup. But while competence is “free” from the perspective of the rulers, resources most certainly are not. Thus, it is always preferable to improve the efficiency of the military by hiring competent officers and reducing its resources than to hire incompetents who may be loyal but who will also squander valuable resources. We provide empirical evidence for our argument by examining the fate of high-ranking officers in Saddam Hussein’s military in the aftermath of the Iran-Iraq War.

1 The Guardianship Dilemma

The Guardianship Dilemma has plagued regimes for centuries, and has proven a tough challenge even during the last century, when militaries deposed more rulers than all other forms of political instability combined. Between 1945 and 2002, more than two-thirds of the extraconstitutional leadership changes in dictatorships were caused by coups d’etat (Svolik, 2009). Even among *all* leaderships changes between 1919 and 2004, the 260 coups represent nearly 9% (Goemans, Gleditsch and Chiozza, 2009).

The Guardianship Dilemma depends on the threat environment faced by states’ political regimes: stronger external threats increase the need for a powerful military, but the more powerful the military, the more dangerous it can be to the regime’s political autonomy and even its survival (Huntington, 1957; Finer, 1988; Acemoglu, Ticchi and Vindigni, 2008; Feaver, 1999; Svolik, 2012).

Some exceptional studies do hold that the strong foreign threats can enhance civilian control of the armed forces. Desch (1999) argues that civil-military relations depend on the nature of the threat environment faced by the state. Civilians have greatest control over the military when the state faces grave international threats, and least control when the state must deal with domestic challenges. Staniland (2008) adds that this relationship depends on the extent to which the regime is deemed legitimate and is adequately institutionalized. It is not, however, clear why one should expect militaries to have the least influence over policy when their services are in highest demand, or why domestic and international threats should have such dramatically different effects on civil-military relations.

Rulers who worry about their own militaries do have another option at their disposal: they can purposefully limit the strength of their armed forces, leaving them too weak to mount a coup but also making them grossly ineffective against the ex-

ternal threat (Svolik, 2013; Feaver, 1996, 154). At its most extreme, this strategy could deprive the state of a military altogether, as it has done in Costa Rica and some remote island states such as Kiribati and Samoa. More realistically though, most rulers must content themselves with finding ways of maintaining reasonably useful guardians without becoming their victims. This is what nearly all studies of civil-military relations investigate as well. Strategies discussed vary from institutional constraints involving limitations on the autonomy of military organizations and the creation of parallel armed forces, to efforts to control the disposition of military agents by providing patronage or by fostering a professional culture among military personnel (Finer, 1988; Pilster and Boehmelt, 2011; Quinlivan, 1999; Powell, 2012).

A key strategy for ensuring the loyalty of militaries is to control the membership of the officer corps. The calculus for rulers in this context is straightforward: where the loyalty of potential guardians might be questionable, appoint those for whom the status-quo is most profitable and who will therefore have the weakest incentives to overthrow the ruler. These privileged groups are generally among the regime's "communities of trust" (Enloe, 1975; Quinlivan, 1999). Saddam Hussein, for instance, pulled heavily from his minority Sunni-Arab ethnoreligious group when selecting personnel for the Iraqi security apparatus (al-Marashi, 2002). Scholars have often argued that making these personnel decisions on the basis of implied loyalty rather than competence can seriously erode the military effectiveness of the armed forces (Huntington, 1957; Brooks, 1998; Gaub, 2013, 231-2). Some have even gone as far as suggesting that rulers might actually recruit less competent officers on purpose as a means of ensuring their loyalty (Egorov and Sonin, 2011). We shall have an occasion to address these particular claims.

While the logic of the Guardianship Dilemma serves as the foundation of our existing understanding of civil-military relations, one must wonder if this logic is convincing. Most of the work that explores this logic explicitly suggests that the problem turns on the ability of the rulership to commit to resource transfers or policy concessions to the military. A failure to manage the dilemma then reflects features of the social, political, or economic environment that prevent rulers from making credible promises to their guardians. Besley and Robinson (2010) argue that if social conflict over public spending is serious enough, rulers cannot commit to paying a wage that is sufficiently high to ensure military loyalty when the armed forces are optimally sized. The ruler's best response is to limit the size of the armed forces and avoid a coup altogether at the cost of having a weaker military. Leaving aside the question of why authoritarian rulers would be unable to secure sufficient resources for the military, the theory cannot explain why coups occur; after all, the equilibrium probability of a coup is zero.

Acemoglu, Ticchi and Vindigni (2008) do tackle this question head on. In their view, the transition from autocracy to democracy can end in a coup because the incoming democratic regime no longer needs a well-paid military for internal re-

pression. Since the democratic government cannot credibly commit not to reform the military after it comes to power, the military has incentives to prevent the transition. It is, however, unclear why the government would not be able to make such a commitment: since the military cannot be disbanded overnight, the continuing threat of a coup should give the government enough reasons to maintain the high wages.

Other studies suggest that the problem arises from political or military elites holding private information about relevant features of the strategic environment. Informational asymmetries are common in civil-military politics, reflecting the functional differentiation in tasks between political and military actors, and are troublesome to the extent that militaries and political regimes have dissimilar preferences over outcomes (Brooks, 2008). Svoblik (2013, 2012) argues that militaries leverage their coercive power to demand favorable policies from the regime. Coups can occur when the military believes that the ruler has reneged on their agreement, which can happen because the military lacks complete information about the government's activities. However, even though some policies might be opaque to the military, most large issues — such as the military's budget or regulations affecting the armed forces — tend to be highly visible and the policies themselves formulated with the active participation of the military.

In this respect, Egorov and Sonin's (2011) assumption that the ruler's agent ("vizier") has private information about the threat environment is much more plausible. A competent agent is more likely to observe whether the enemy is weak, and so its incentive to betray the ruler by doing nothing to counter that threat is higher. To counter this, the ruler hires a less competent agent and since the expected loyalty is higher, the required pay is lower. The only reason the ruler does not hire total dolts is because their inability to distinguish whether the enemy is weak would cause them to squander valuable resources. One might wonder about a notion of competence that is unrelated to the agent's ability to defeat the enemy. After all, if the enemy is more likely to prevail in the presence of an incompetent vizier, the ruler's money-saving imperative that drives down his desire to hire a competent agent will be, well, much less imperative.

In order to assess the logic of the Guardianship Dilemma, we distill the dilemma to the most essential characteristics identified by previous research. (1) The leaders of political regimes must defend against external threats. Unfortunately for them, the guardians appointed to defend the state can also be a threat to the regime. (2) Rulers have the 'power of the purse', and manage the flow of resources to the armed forces in response to both external threats and the risk of a coup. (3) Rulers control who is charged with running the state's armed forces, and may select these agents on the basis of both competence and their affiliation with a social, political, or economic group. (4) The more competent the military agents, the more likely they are to prevail against the external threat *and* against the ruler should they choose to execute a coup.

2 The Model

Consider a model with two players, R (ruler of the political regime) and G (general).¹ The status quo distribution of benefits in this society privileges certain groups over others, and may be based on ethnicity, religion, geography, or other cleavages within the state. In South Africa during Apartheid, for example, a racial divide between whites and non-whites determined access to social, political, and economic opportunities (Thompson, 2001). For the sake of parsimony, we abstract away from the precise nature of these cleavages and assume simply that the benefit a member of some group i derives from the status quo is $b_i \in [\underline{b}, \bar{b}]$ such that $0 < \underline{b} < \bar{b} < 1$. Some groups have higher status quo benefits than others, so their incentive to overthrow the regime will be weaker. Ugandan ruler Milton Obote was a Northerner, and knew that for fellow Northerners, a coup to address ethnic grievances would be unnecessary. The loyalty of the alienated Southerners, in contrast, was far more questionable (Horowitz, 1985, 488, 501). These status quo benefits are normalized such that 0 represents obtaining nothing (e.g., because one is dead or in prison) and 1 represents the maximal benefit of personal rule. We shall further normalize R 's competence to 1, and her security resources to 1.

The timing of the game is as follows. The ruler chooses the group from which to pick the general, b_i , his level of competence, $\theta \in [0, \bar{\theta}]$, and the amount of military resources to make available to him, $m \geq 0$. The marginal cost of giving the general a unit of military resource is 1. All these parameters are observable by the ruler when making her choices, and known to the general selected.

Powerful, well-endowed military forces are more likely to succeed in battle against an external threat. However, if the military attempts a coup, strong forces are also more likely to overcome the defenses that protect the regime, such as paramilitary units and pro-government militias, and can more easily capture strategic targets and members of the government (Powell, 2012, 1024). Furthermore, the ability of competent military leaders to marshal forces effectively is vital when the government faces external threats, but can be particularly dangerous if cunning generals turn against the regime. Returning again to the case of Uganda, Milton Obote began to fear General Idi Amin's wiles. When the threat posed by Amin to the regime became clear enough, Obote demoted the general in an attempt to limit his influence (Horowitz, 1985).

We represent the probability that an actor of competence θ in control of military resources m_1 prevails against an opponent — here, either the external threat or the ruler's own defenses — with resources m_2 with the familiar ratio contest-success function:

$$p(m_1, m_2; \theta) = \frac{\theta m_1}{\theta m_1 + m_2}.$$

¹For simplicity, we shall refer to the ruler as “she” and the general as “he”.

Following the ruler’s choice of a general and military resources, the selected general decides whether to execute a coup or remain loyal. If he executes a coup, he takes over with probability $p(m, 1; \theta)$, in which case his eventual benefit goes to 1, and he is defeated with complementary probability, in which case his benefit goes to 0. The coup is costly for the general: $c > 0$.

After the coup decision, the external threat of size $T > 0$ is realized. It is important to note that by “external” threat, we mean any threat from outside of the government — whether foreign or domestic — that threatens the survival of the rulership.² If the general is still around when this threat is faced, (because he remained loyal or after a successful coup), he defeats this threat with probability $p(m, T; \theta)$, in which case he obtains his benefit (b_i if he was loyal to the ruler and 1 if he took over in a coup), and he is defeated with complementary probability, in which case his benefit is 0. If the general was removed after an attempted coup, the ruler herself faces the external threat, and defeats it with probability $p(m, T; 1)$, in which case she retains power with a benefit of 1, and is herself defeated with complementary probability, in which case her benefit is 0.

The highest expected payoff from a coup in the absence of an external threat is when there is no risk: $1 - c$. If G would not want to execute a coup even when a favorable outcome is certain, then he will always remain loyal irrespective of his competence and resources. To make the model interesting, we shall assume that G ’s loyalty is not so easily ensured:

ASSUMPTION 1. Every general is a would-be ruler: $\bar{b} + c < 1$.

We also make several additional assumptions because we want to focus on the basic tension between security against external threats and security against the force that is supposed to defend against these threats. Some are made for technical convenience and have no bearing on the results, while others can be defended on empirical grounds.

First, the benefit of membership in some group, b_i , is fixed and not borne by the ruler. We can think of this parameter as the consequence of social, economic, and legal institutions underpinning the order in the polity, and as such not really available to the ruler for private consumption. It is the “cost of doing business” and it is not the case that if the ruler picks a general from a less advantaged stratum then he would be saving on that cost.

Second, the ruler pays the cost of resources she transfers to the general but there is no budget constraint. It is highly unlikely that any particular general would be so expensive to get as to trigger a budget constraint, or that rulers are particularly constrained by budgets when it comes to their desire to endow the military with

²We adopt a broad conception of the external threat in order to understand the fundamental dynamic illuminated by the model. Once we have done this, we can ask how international and domestic threats might differ in relevant ways.

resources. If we were to assume that there was a hard budget constraint, then we might find that the ruler limits the size of the military because of poverty, and not because of any security issues, which is the goal of our analysis. As a result, even if one wanted to introduce a budget constraint, one would have to argue that it would bind, and even when it does bind all it will do is introduce a cap on military allocations, suppressing the mechanism we have in mind. We do assume that these transfers are costly, so there is a disincentive to put too much into the military, but as we shall see this is not going to be the concern in general.

Third, the ruler pays no cost if a coup occurs. Alternatively, the ruler can be assumed to pay a cost if a coup occurs without altering anything in the analysis except carrying another parameter across all calculations. The fact that a coup can depose her with positive probability is already an inducement for her to want to avoid it. Adding extra costs simply strengthens an incentive that is already present (and, as we shall see, quite strong). We do want to consider costs for the general, however, because he has to make a decision about executing a coup, and because these costs might represent institutional features of the existing regime that need to be taken into account.

Fourth, neither the ruler nor the general pay any costs when they fight the external threat. This assumption is consistent with the structure of the model, which allows for no choice to avoid that threat: the ruler simply must deal with it. Since the size of the threat already allows us to capture just how bad it could be for her, there is no need to introduce additional costs.³

Fifth, both the ruler and the general are constrained in that they can only use the available military forces m to deal with the external threat (in effect, the ruler's internal security designed to deal with coup attempts is not useful against external threats). This might appear too restrictive because it seems to disregard strategies the ruler could use to decrease the incentives for a coup. Such coup-proofing tactics could involve increasing the resources devoted to internal security, making it difficult for military units to coordinate and communicate without passing through centralized channels, dispersing units or staffing them with non-locals, and others. These measures would decrease the probability of coup success at any level of military resources and increase the costs of launching one.

While our cost parameter can capture some aspects of these tactics, it cannot capture others. We could model the effect of such tactics on the probability of coup success with the resources available to the ruler's internal security forces or her competence, but we have normalized them both. The results will not change if we were to use variables instead as long as we take them as given. In other words, if we think of the Guardianship Dilemma in the context of the ruler having done

³Moreover, if we were to introduce costs of fighting, we have to be careful with the general's payoff if he is eliminated in a coup. Since he would face no enemy when the coup attempt fails, if the costs of fighting the enemy are sufficiently high, he would execute a coup simply because of the chance of failure that would allow him to evade paying these costs.

everything possible to minimize the internal dangers, our analysis follows without any changes. The only downside is that one could not take the model to data where these coup-proofing measures vary without some straightforward modifications that will not alter any of the substance of our argument.

This assumption might also be criticized on grounds that internal security forces may improve the the state's ability to defeat an external threat by augmenting the might of the regular military. We could account for this possibility by incorporating the ruler's own internal security resources in the probability of success against T . Because this merely involves adding a constant, however, our analysis will not change. We are also doubtful about whether this addition would be appropriate. For instance, one might follow Svolik (2012), who makes an empirical claim that the army is only generally useful for dealing with mass revolts or foreign forces. In most states — especially dictatorships — the day-to-day security is managed by another apparatus, whose personnel are generally not useful for large-scale operations. One could also point out that some coup-proofing measures (e.g., making it difficult for the commanders to coordinate) might actually have a detrimental effect on the state's ability to defeat the external threat.

Sixth, the resources given to the military are equally useful for a coup and for fighting the external threat. One might question this on two grounds: it could be that resources are not fungible, and even if they are, they might be useful only in one of the two situations. For example, salaries, health and pension benefits, and payments to civilian contractors are certainly included in military budgets, but they are not likely to increase the fighting ability directly. Spending on some types of technologies could improve the fighting ability when it comes to the external threat without being very useful in a coup. Submarines and fighter jets might belong to that category.

Although both points are doubtless correct, they have only tangential bearing on the Guardianship Dilemma. Since the first type of spending will affect the incentive to launch a coup through the benefits derived from the existing regime, its effects can be approximated by the benefit parameter b_i . The only difference, of course, is that since these payments are part of the budget, they would be costly to the ruler. Under the no budget constraint assumption, this would merely result in another parameter being subtracted from the ruler's payoff, necessitating further assumptions about the marginal costs of these funds, and perhaps restrictions that ensure an interior solution. That solution, however, is not going to produce any difference in the dynamics we study.

Finally, one might wonder about the assumption that more resources given to the military must necessarily increase its ability to prevail in a coup, as built into the functional form of $p(\cdot)$. From the perspective of any coup-plotters, there are some fundamental problems that they would need to overcome before having any chance of success: collective action and coordination problems that arise from incentives to renege on the plot and the necessity of conducting preparations in secrecy, as

well as the absolutely critical question about securing the cooperation or at least neutralizing the units in the armed forces that were not privy to the coup (Luttwak, 1979). Very plausibly, these problems, and especially the latter, might be quite aggravated by the size and complexity of the armed forces. In other words, it could be that beyond some level, the larger and more organizationally complex the military, the harder it is for any general to organize a coup. (Such a dynamic could account for the political quiescence of the armies of the Soviet Union, North Korea, and China.) While this is certainly an intriguing possibility, we believe that such a “pacifying” dynamic would have to be quite exceptional as most countries do not have the population base to maintain large armies. Note also that an attempt to construct and maintain an army of sufficient size would probably run afoul of resource constraints.

In the end, ours is emphatically not a general model of coups. We do not study how coups are organized and how they succeed (Sutter, 2000). Our interest is in the fundamental Guardianship Dilemma, which has to be analyzed prior to dealing with any strategies for ameliorating its effects. To this end, we have stripped the model of any factors that are not essential to the dilemma, and whose presence might obscure rather than clarify its logic.

3 Known External Threat

We shall begin our analysis with the case where $T > 0$ is common knowledge. If R does not hire a general, then there is no threat of a coup, so R 's payoff is $p(m, T; 1) - m$; i.e., she simply has to meet the threat with her own competence and the resources she has allocated. Maximizing this payoff yields $\tilde{m} = \max(0, \sqrt{T} - T)$. In any equilibrium in which G gets hired, R 's expected payoff must exceed the baseline of $p(\tilde{m}, T; 1) - \tilde{m}$. This immediately implies that no coup can occur in equilibrium. In such an equilibrium, R 's expected payoff would be

$$p(1, \theta m; 1) p(m, T; 1) - m < p(m, T; 1) - m \leq p(\tilde{m}, T; 1) - \tilde{m},$$

which means that she would strictly prefer not to hire a general in the first place. The only equilibrium possibilities, then, are that either no coup occurs at all or one occurs with positive probability less than one. The following result (all proofs are in Appendix A, which is available online) shows, among other things, that a coup can never occur in equilibrium.

LEMMA 1. *In any equilibrium, G remains loyal if, and only if, $T \geq T_i^*(m, \theta)$, where*

$$T_i^*(m, \theta) = \left(\frac{\theta m}{c} \right) \left[\frac{\theta m}{\theta m + 1} - (b_i + c) \right], \quad (1)$$

with T_i^ increasing in both parameters whenever it is non-negative.* \square

This immediately tells us that if G would remain loyal in the absence of an external threat, then he will remain loyal in the presence of such a threat irrespective of its size. If, on the other hand, G would be disloyal in the absence of an external threat, then he will execute a coup in the presence of an external threat only if this threat is not too large $T < T_i^*(m, \theta)$. In this sense, sufficiently grave external threats can discipline even a potentially disloyal general and deter him from executing a coup, a sort of “circling the wagons” effect. This effect is due to the fact that the general only wants to take the risks and pay the costs of a coup when he is sufficiently confident about surviving the conflict with T , since survival is necessary to reap the benefits of ruling the state. In turn, as T increases, the loyalty-inducing effect of this external threat allows the ruler to pour additional resources into the military without triggering a coup.

Moreover, since R would not hire a general if a coup is certain, and Lemma 1 shows that G must remain loyal when indifferent, it follows that in equilibrium the probability of a coup must be zero. This leads to the following result.

LEMMA 2. *Fix any social group b_i . If R 's choices ensure G 's loyalty, then R always picks the most competent general from this group, $\bar{\theta}$, and endows him with:*

$$m_i^*(T) = \begin{cases} \max\left(0, \sqrt{T/\bar{\theta}} - T/\bar{\theta}\right) & \text{if } S(T) \leq S_i^*(T) \\ S_i^*(T)/\bar{\theta} & \text{otherwise,} \end{cases}$$

where

$$S_i^*(T) = \frac{b_i + c + cT + \sqrt{(b_i + c - cT)^2 + 4cT}}{2(1 - (b_i + c))} \quad (2)$$

is the maximum level of disloyalty that would not provoke a coup, and

$$S(T) = \sqrt{\bar{\theta}T} - T \quad (3)$$

is the level of disloyalty for the most competent G with resources optimally provided to deal with the external threat. \square

This tells us how R would allocate military resources if doing so would preserve the loyalty of the general. We now show when R would prefer to hire a general given that she would have to ensure his loyalty.

LEMMA 3. *In equilibrium,*

- (i) R never hires G if the maximum competence is worse than her own: $\bar{\theta} < 1$;
- (ii) R always hires G with $\bar{\theta} > 1$ when the external threat is sufficiently large: $T \geq 1$;

(iii) *R may or may not hire G with $\bar{\theta} > 1$ when $T < 1$, depending on the costs of a coup (c) and the benefits from the status quo (b_i). In particular, if both are sufficiently small, then R will not hire anyone.* \square

We have now established that the ruler will never hire anyone less competent than herself and that whenever she chooses to hire a general, she picks the most competent one she can find. Moreover, if the external threat is sufficiently serious, the ruler always hires a general although she might have to ensure his loyalty by providing him with fewer resources than what is optimal for dealing with that threat.

Cases (i) and (iii) of Lemma 3 are substantively unlikely. The former essentially means that no potential general is more competent than the ruler, a highly unlikely scenario (well, except perhaps if the ruler is Napoleon, but even then there might be a potential Wellington!). The latter requires that the external threat be negligible, in which case it is very easy to trigger the disloyalty of any general, which is why the decision to hire depends only on the benefits of the status quo and the costs of a coup. Since the ruler's incentive to hire a general turns on a looming external threat and the need to get someone competent to deal with it, this case is irrelevant for our purposes. Consequently, we shall exclude these substantively unappealing scenarios from further consideration:

ASSUMPTION 2 (PREFERENCE FOR HIRING). There always exist generals more competent than the ruler ($\bar{\theta} > 1$), and the external threat is never negligible ($T \geq 1$).

Under Assumption 2, Lemma 3 implies that R will always hire a general in equilibrium. The following result shows that, generally speaking, the ruler will give preference to the privileged groups when it comes to selecting a general.

LEMMA 4. *Let b^* be the unique solution to $S_i^*(T) = S(T)$. If $\bar{b} \leq b^*$, then R strictly prefers to pick G from \bar{b} ; otherwise R is indifferent among any $b_i \in (b^*, \bar{b}]$, and strictly prefers any of them to any $b_i < b^*$.* \square

Thus, R will either choose from the most privileged group or from among the few most privileged (when each of them provides enough benefits to ensure the loyalty of generals drawn from them). Moreover, R will always pick the most competent G she can although she might have to handicap the general resource-wise in order to ensure his loyalty. We can state the main result somewhat loosely as follows.

PROPOSITION 1. *If the extent of the external threat is common knowledge and assumptions 1 and 2 are satisfied, then in any subgame-perfect equilibrium, the ruler picks the most competent general. If there are groups that derive sufficient benefits from the status quo to ensure the loyalty of a general selected from them at the allocation that is militarily optimal to deal with the external threat, then the ruler chooses from any among them, and endows the general with the optimal resources*

(the equilibria are payoff-equivalent). If no such group exists, the ruler selects the general from the most privileged group, and endows him with just enough resources to ensure his loyalty. No coups occur in equilibrium, but the external threat is not properly met when the ruler is forced to handicap the general. \square

We have thus established that when the size of the threat is known, the Guardianship Dilemma is, in principle, solvable: militaries remain loyal in equilibrium, and the ruler's strategy always privileges domestic political survival over dealing with the external threat. The government hires competent generals, but controls resource flows to the armed forces in order to ensure military loyalty. The more privileged the groups from which the generals are selected, the less biting the trade-off between stability (risk of a coup) and security (risk from the external threat). Since the costs of the coup act as a substitute for benefits, the more effective anti-coup measures, the less biting the trade-off becomes and the less pressing the need to privilege the military. In this way, "coup-proofing" works much as previous studies suggest.

The very solvability of the dilemma and especially the fact that whenever the trade-off between stability and security exists it is always resolved in favor of stability are puzzling given the frequency of military interventions in politics. If rulers have levers for controlling their armed forces, why are defections by military forces such a regular occurrence?

4 Asymmetric Information about the External Threat

Let us now assume that only G observes the actual external threat T , whereas R is only imperfectly informed about it. As before, subgame perfection implies that given an allocation m , a general of competence θ who obtains status quo benefits b_i will execute a coup if, and only if, (2) is not satisfied; i.e., if the threat T is not sufficiently large to deter him. This suggests that it will be sufficient to analyze the case with two types of threats: small and large, with $1 < T_S < T_L$ (notice that we are maintaining Assumption 2). The ruler believes that the threat is T_S with probability $q \in (1, 0)$ and T_L with complementary probability.

From the comparative statics on T , we know that when threats become sufficiently large, the marginal costs of military allocations begin to outweigh their usefulness, so R responds by decreasing m even though there is no danger of a coup. We consider it highly implausible that a ruler will be so hampered by these marginal costs that she would respond to more serious threats by reducing her spending on security. Instead of introducing a parameter for marginal costs and requiring it to be sufficiently small given the maximum threat magnitude, we shall simply restrict the threat to ensure that the optimal allocation is strictly increasing in its size. This is already true when R constrains G , so this really only affects the unconstrained allocation.

ASSUMPTION 3 (REASONABLE COSTS OF SECURITY). The marginal costs of security are not so high as to cause larger threats to require smaller counter-measures under complete information: $m_i^*(T_L) > m_i^*(T_S) > 0$.

If R does not hire a general, her expected payoff is

$$U_A = m \left(\frac{q}{m + T_S} + \frac{1 - q}{m + T_L} \right) - m,$$

which has a unique optimal allocation that results in a strictly positive payoff.

When R hires a general, any allocation can result in one of three outcomes: a certain coup, no coup, and a coup only if the threat is small. To see this, fix some m and observe that if G stays loyal under T_S given that allocation, he must certainly do so under T_L as well. Conversely, if he executes a coup under T_L , then he must also do so under T_S as well. The sole remaining possibility is that he executes a coup under T_S but remains loyal under T_L .

We begin by ruling out the possibility that the ruler will hire anyone when she believes that doing so would result in an inevitable coup (this parallels the complete-information case).

LEMMA 5. *There is no equilibrium in which R hires G when she expects a coup to occur with certainty.* \square

Thus, in any equilibrium in which R hires a general, the general's loyalty is either certain or else only in doubt conditional on the actual size of the threat. The following result shows two things. First, the ruler will never hire anyone less competent than herself. Second, the ruler's strategy depends on her prior belief about the magnitude of the threat. If she is sufficiently convinced that the threat is large (i.e., q is small), then she allocates more resources to G even though she knows that G will execute a coup if the threat is, in fact, small. The allocation is not, however, optimal for meeting T_L either because the possibility that it will be used in a coup against her forces the ruler to curtail it a bit. In this situation, the ruler faces a positive probability of a coup and does not have enough forces to deal with the large threat. If R is sufficiently convinced that the threat is small (i.e., q is high), then she plays it safe: she allocates just enough resources to ensure the loyalty of G under the assumption that the threat is small. While this does ensure that no coup takes place, the ruler will find herself severely handicapped if the threat turns out to be large.

LEMMA 6. *Fix a social group b_i and a level of competence θ . In any equilibrium, R hires G only if $\theta > \max(1, T_S)$. In the unique equilibrium in which R hires G , there exists a unique $q^* \in (0, 1)$ such that*

- if $q \leq q^*$, then R allocates $\min(m_C(q), m_i^*(T_L))$, where $m_C(q)$ is the unique unconstrained maximizer of R 's expected payoff, and G executes a coup if the threat is T_S but remains loyal otherwise (risky strategy);

- if $q > q^*$, then R allocates $m_i^*(T_S)$, and G remains loyal (safe strategy).

When R plays the risky strategy, she not only faces a positive probability of a coup from a general with substantial resources, but may also fail to provide adequate resources to deal with the large threat. When R plays the safe strategy, she certainly fails to provide adequate resources for the large threat. \square

Having established what resources R will allocate once she has chosen G with some competence θ from some class b_i , we now ask how she makes these selections. Since Lemma 6 shows that hiring can only occur if $\theta > \max(1, T_S)$, we shall assume that $\bar{\theta}$ satisfies this condition.

LEMMA 7. *Fix a social group b_i . In any equilibrium, R hires the most competent G she can ($\bar{\theta}$).* \square

Finally, we need to consider the social group from which R selects the general. We first show that R 's payoff is non-decreasing in b_i if she pursues the riskless strategy. In particular, it is constant in b_i if the complete-allocation optimum against T_S is unconstrained, and strictly increasing otherwise. Thus, starting with a very low b_i the payoff will not change, and increasing b_i eventually causes it to start increasing.

LEMMA 8. *Let $b^*(T)$ denote the unique solution to $S(T) = S_i^*(T)$. And let $b_1 = \min(b^*(T_S), b^*(T_L))$ and $b_2 = \max(b^*(T_S), b^*(T_L))$. If $\bar{b} \leq b_1$, then R strictly prefers to pick G from \bar{b} . If $\bar{b} \in (b_1, b_2)$, then R strictly prefers to pick G from \bar{b} for $q > q^*$ if $b_1 = b^*(T_S)$, and for $q \leq q^*$ if $b_1 = b^*(T_L)$, and is indifferent among any $b_i \in [b_1, \bar{b}]$ for any other q (but strictly prefers any of them to $b_i < b_1$). If $\bar{b} \geq b_2$, then R is indifferent among any $b_i \in [b_2, \bar{b}]$ (but strictly prefers any of them to $b_i < b_2$).* \square

We can now state the main result under asymmetric information.

PROPOSITION 2. *If only the general knows the extent of the external threat, then in the essentially unique equilibrium the ruler picks the most competent general from the most privileged strata in society. If the ruler is sufficiently sure that the threat is small, she provides the general with only enough resources to meet that threat (even these might be constrained), and the general remains loyal regardless of the extent of the threat. If the ruler is sufficiently sure that the threat is large, she provides the general with resources that balance the risk of a coup with the risk of failing to meet the large threat with adequate resources (even these might be insufficient for the large threat). The general remains loyal if the threat is large but executes a coup if the threat is small.* \square

Proof. The result follows immediately from lemmata 6, 7, and 8. The equilibrium is essentially unique because R might be indifferent among many values of b_i as long as they are sufficiently high. Each of these corresponds to a different equilibrium but they are all payoff-equivalent. \blacksquare

5 Discussion

5.1 External Threats and Military (Dis)Loyalty

Although it appears to make perfect sense, this Guardianship Dilemma turns out to be incomplete. It begins with the premise that the threat environment will create the need for armed forces, which in turn will pose yet another risk for the regime, but fails to consider what effect this environment will have on that new risk. At best, the Guardianship Dilemma offers a straightforward linear extrapolation: the worse the threat environment, the greater the need for armed forces, and, if this need is met, the larger the risk they will pose.

What is missing in this logical chain, however, is the simple fact that if the military does execute a coup and take over the government, the original threat is not going to magically disappear. The new rulers will have to face many, if not all, of the same problems and dangers that had confronted the old ones. The Malian regime of Amadou Toumani Toure, for example, was overthrown in a military coup d'état in March 2012. Even though the regime had been deposed by the military, the state was still forced to deal with an ongoing rebellion by Tuareg fighters (Nossiter, 2012). Similarly, Syria experienced no fewer than eight successful coups d'état between 1950 and 1970, when Hafez al-Assad assumed power (Pipes, 1989; Powell and Thyne, 2010). Despite the frequent changes in rulership during this period, relations between the Arab state and its primary opponent, Israel, remained tense (Neff, 1994).

The persistence of threats across regimes is a very real and important consideration for military agents who are considering whether or not to intervene in politics. Because these forces must both overthrow the regime *and* face the threat, external foes help to induce loyalty by a state's military forces. This "circling of the wagons" effect is shown in Figure 1a, where we focus on threats that are at least moderately large ($T > 1$).⁴ If the external threat is grave enough ($T \geq T_i^*(m, \theta)$), rulers can devote the optimal allocation to defense without triggering a coup. In the case of Iraq, President Saddam Hussein was able to relax constraints on the Iraqi military during the Iran-Iraq War principally because these forces were fighting for the survival of the state (Hiro, 1991; Pelletiere and Johnson, 1991). However, rulers in this context must also defeat a stronger threat, which discounts the probability of survival (see Figure 1b).

Alternatively, allocating the optimal amount of resources for defense would trigger a coup when $T < T_i^*(m, \theta)$, leading to a strictly lower chance of survival for the rulership. In this case, the ruler is safer by *reducing* the amount of resources that she devotes to defense to the coup-constrained amount, even though this will handicap the military. Muammar Qaddafi of Libya, for example, led a regime that

⁴The parameters for all plots are: $b = 0.25$, $c = 0.30$, $\theta = 16$, $T_S = 1$, and $T_L = 4$.


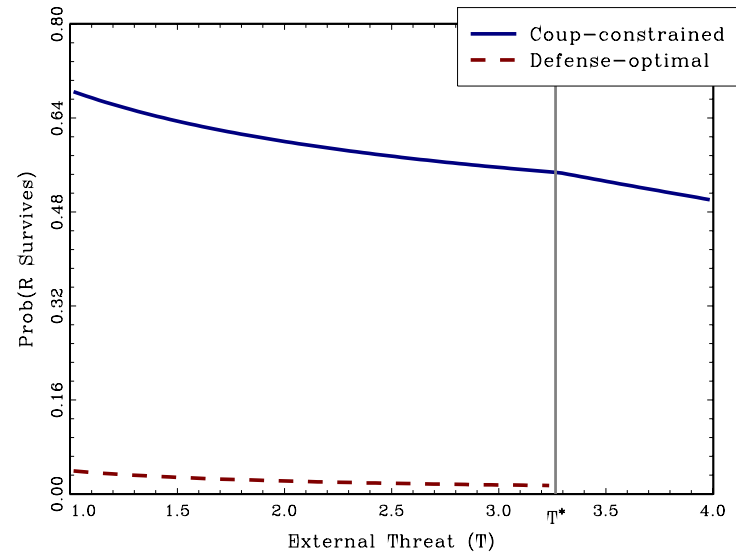
(a) R 's allocation to G (b) R 's probability of survival

Figure 1: Military Endowment and Regime Stability with Known Threats

faced only moderate levels of external threat for most of his 42 year rule. When it came to managing the military, Qaddafi — who himself had taken power in a coup — purposefully limited its power in order to improve regime security (Lutterbeck, 2013, 40).

By identifying circumstances under which rulers will withhold resources from their militaries, our study builds on the work of Besley and Robinson (2010), who demonstrate that rulers will sometimes keep their militaries smaller than optimal in order to ensure their loyalty. The mechanisms driving these constraints, however, are quite different. Constraints in the Besley and Robinson (2010) model are a consequence of social conflict over public spending, which prevents the regime from credibly committing to resource transfers. Because the military's loyalty cannot be purchased under such conditions, constraints on the strength of the armed forces are necessary to prevent defection. In contrast, our model shows that constraints can be a function of the threat environment even when rulers can make credible commitments (in our model the resources are given before the coup choice). Rulers can leave the military unconstrained so long as the external threat is sufficiently large, but must impose limitations on their armed forces when faced with intermediate threats.

In one way, these theoretical results are consistent with the basic claim of Desch (1999): militaries are less willing to intervene when external threats loom large. However, our model reveals that with a known external threat, the loyalty of the military does not depend on the size of that threat. The ruler can always remain safe from a coup, whether the external threat is small or large, by controlling the power of the armed forces. This suggests that the mechanisms studied by Desch cannot account for the dilemma despite identifying the loyalty-inducing effects of larger threats. The risk of military disloyalty is not due to variation in the threat environment, but is instead triggered by another factor, one that prevents rulers from calculating and appropriating the correct level of military resources for the given environment. This factor is the asymmetric information that the military and the ruler might have about the seriousness of the threat.

5.2 Disagreements about the Threat Environment

Delegating the responsibility for defense to the armed forces creates a less obvious, but, in many ways, more vexing problem for rulers. Militaries are maintained because they possess specialized skills and tools for assessing and combating the state's enemies. This specialization means that militaries will possess private information about the nature of the threat environment — information that we show is key for rulers who are trying to navigate the Guardianship Dilemma. Rulers' beliefs about the threat environment determine the amount of resources they devote to the military, which, in turn, drives both the risk of a coup and the ability to defeat external threats.

Figure 2 shows that when the ruler believes the threat to be small ($q > q^*$), she allocates a relatively low level of resources to the military, which ensures its loyalty whether the threat is actually small or large. However, if the ruler mistakenly underestimates the threat, she will allocate too few resources for facing the large threat. While the ruler remains safe from a coup, the regime is more vulnerable to external foes. This danger is illustrated in 3, which shows the probability of regime survival as a function of the ruler's belief.

Alternatively, when the ruler believes the threat to be large ($q \leq q^*$), she allocates a level of resources that will trigger a coup if the threat is actually small, but helps the regime defend against large threats. Because rulers are strictly worse off in a coup, the risk of overestimating the threat is the most dangerous possibility facing their regimes. As a result, even at this high tier of resource allocation, the ruler hedges against the risk of a coup by imposing slight constraints on the military unless she is absolutely certain that the threat is large. So when the ruler must deal with uncertainty, she faces the possibility of either overestimating the threat and risking a coup, or underestimating it and leaving the state exposed to enemies. Rulers are safer when responding appropriately to the given threat environment, and safest when facing a definite small threat.

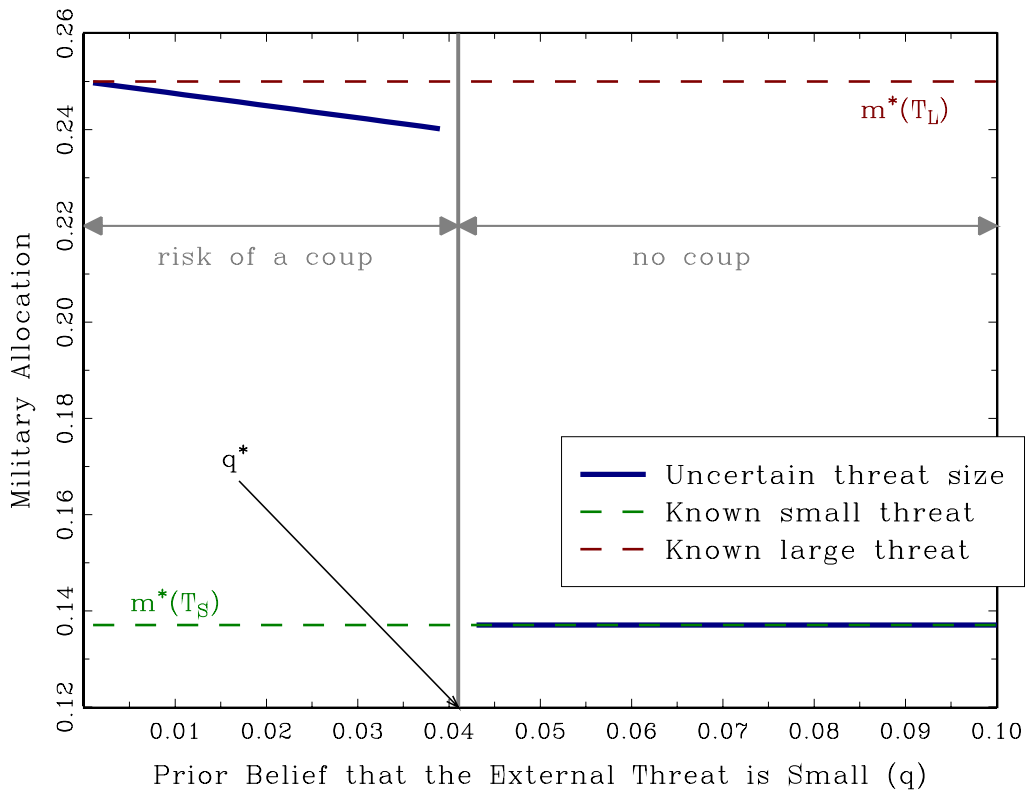


Figure 2: Threat Estimates and Resource Allocation




Figure 3: The Divergence of Threat Estimates and Regime Security

These findings complement the work of researchers like Svobik (2012) and Brooks (2008), who posit that informational asymmetries can complicate the civil-military relations within states as well as the response of regimes to external threats. While Svobik (2012, 2013) focuses on the case in which military agents are asymmetrically informed about the policies instituted by a regime, we outline the difficulties that arise from the private information that militaries gain while fulfilling their responsibilities as guardians of the state. In this way, we characterize the essence of the dilemma inherent in civil-military relations: the competencies that make military agents effective also make them a threat.⁵

5.3 Relating the Model to Empirical Work

A small, but growing, body of scholarship analyzes correlations between coups and variables that could be interpreted in terms of our model. We now discuss how one might account for these empirical regularities with the mechanism it identifies, while also highlighting the additional assumptions one would need to make and the

⁵We discuss why the military may withhold information about the threat environment in Appendix C, available online.

difficulties with positing some of the direct relationships suggested by these studies.

The model predicts that external threats help to induce military loyalty, though civilian control ultimately depends on whether or not the ruler possesses accurate information about the threat environment. In this context, it is useful to consider the empirical results of two studies that find that the probability of coups is lower if the country is involved in a war (Talmadge and Piplani, 2014) and even in a crisis (Arbatli and Arbatli, 2014). The explanatory mechanisms these studies offer are different (although not necessarily incompatible): Talmadge and Piplani (2014) argue that when the military is engaged in a war there are fewer opportunities for a coup and more uncertainty about who will join it, whereas Arbatli and Arbatli (2014) argue that crises allow rulers to commit credibly to transfers to the military and to generate rally-around-the-flag effects. Even though either one or both of these mechanisms could be relevant, it is worth noting that our model could produce these predictions in a very straightforward manner.

Since the key variable is the extent of disagreement about the severity of the threat, which is difficult to measure directly, one might wish to conceptualize the uncertainty about the threat in terms of factors that make it more or less likely for such disagreements to arise. For instance, an ongoing war would be indicative of a fairly serious threat that neither the military nor the ruler could possibly be in doubt about. Moreover, since longer wars can potentially reveal more information (Slantchev, 2003), the longer the war, the less likely disagreement should be. Analogously, a crisis could indicate a somewhat less severe threat with some possible disagreements because of diverging estimates about the likely outcome of the crisis. When the country is at peace and not involved in a crisis, on the other hand, there is no clear evidence that could force the political and military estimates to converge: since all threats are purely hypothetical at that point, the possibilities for different opinions relying on different pieces of information would proliferate. The longer the peace spell, the more likely are these differences to become serious disagreements.

In other words, one might think of disagreement as a continuous variable, proxied by how long the country has been at peace, whether it is involved in a crisis, and whether it is actively fighting. The model would then predict that coups are most likely when there is peace (and the longer the peace spell, the higher the probability of a coup), significantly less likely when the country is involved in a crisis, and quite unlikely when it is involved in a war (and the longer the war, the lower the probability of a coup). Thus, our model can account for the correlations found by both of these analyses without having to resort to different explanatory mechanisms.

While uncertainty about the severity of the threat has a straightforward direct effect, its role as a mediator for the effects of other variables is more complex. Consider, for instance, the problem of relating the military's endowment to the probability of a coup, which is the subject of a study by Powell (2012). Using an expected utility framework, Powell (2012, 1021) notes that the military would

be more likely to execute a coup if it anticipates high benefits from doing so, and if it believes that it has a high probability of success. He then argues that higher (or increasing) levels of funding per soldier will lower the probability of coups (Hypothesis 1) but make coups more likely to succeed (Hypothesis 5), which he considers a paradox (Powell, 2012, 1025).

Relating these two hypotheses to variables in our model is not as easy as one might think because of the way they treat military resources. The model's basic assumption is that Hypothesis 5 is correct: this is built into the functional form of the probability of success, which increases in the amount of resources controlled by the military. Furthermore, the model also assumes that these resources cannot be used to increase directly the benefit of a coup, as the reasoning behind Hypothesis 1 would have it. On the other hand, the model does allow the ruler to select generals from more privileged groups, which decreases the potential gains from a coup, but does not allow this to influence the probability of success. We think that there are substantial analytical benefits to be had from keeping these effects separate. After all, even if resources are in the infinitely fungible medium of money, it is not a simple matter to explain how they could create the supposed paradox: higher salaries would not necessarily translate into better training, while purchasing better military equipment would not necessarily yield the military higher benefits from the status quo.

Since Powell's (2012) statistical analysis uses the government's military expenditures, which include everything from salaries to equipment — not to mention cases where a significant chunk of the military budget goes to civilian contractors and employees, which increases neither the status quo benefits of the military nor its ability to prevail in a coup — to measure the resources made available to the military, it is not possible to map the findings to the model. Unfortunately, neither it is possible to draw the conclusions he does from that analysis. For instance, one of the findings is that contrary to Hypothesis 5, larger expenditures do not increase the probability of coup success. This would be just as stunning if the spending was on better training, better equipment, and better organization, as it would be trivial if it was on better salaries, health and pension benefits, and other perks. In other words, without disaggregating military expenditures to distinguish between spending that could potentially improve the capabilities from spending that is designed to improve the status quo benefits of the military, one cannot take this finding as contradicting the assumption of our model that resources designed to improve capabilities would increase the probability of success. Instead, the model's clear conceptual distinction highlights a vague and under-theorized aspect of the explanatory mechanism that generates Hypothesis 5.

It is perhaps even more interesting to attempt to relate Hypothesis 1 to the model, at least when it comes to the causal mechanism. (Obviously we cannot make much of the correlation findings here either for the reasons discussed above.) Powell's (2012) hypothesis is that better-endowed militaries (or those that enjoy an increase

in resources) should be less likely to stage a coup. Since under complete information about the threat the probability of a coup is constant at zero in equilibrium, any variation has to come from the asymmetric information case, where the relationship between resources and coup probability is mediated by the extent of disagreement about that threat.

On one hand, Figure 2 seems to predict precisely the opposite relationship to the one stated in Hypothesis 1: the only positive risk of a coup happens when the ruler is sufficiently convinced that the threat is large and so provides the military with a lot of resources ($q < q^*$). Since in all other cases the ruler opts for the small allocation and no risk, one might be tempted to conclude that militaries with more resources are more likely to execute a coup. However, if we consider the dynamics in the range of the parameter space where a risk of a coup exists, we find something different. As the plot shows, the military's allocation is increasing in the ruler's belief that the threat is large. Recall now that in equilibrium the military executes a coup only when the threat is actually small and assume that the ruler is not deluded on average (meaning that as her belief that the threat is large goes up, the actual probability that it is small is also going down). In this situation, we would expect the probability of a coup to be decreasing as the resources increase. Overall, the model would lead us to expect that militaries with fewer resources do not generally engage in coups, but also that when coups do occur, militaries with more resources are less likely to have caused them. While the latter is consistent with Powell's (2012) findings, we should not read this as some sort of unequivocal support for the model: after all, our mechanism does depend on the crucial intervening variable of the degree of disagreement, and this is naturally absent in the estimations that take a completely different mechanism as their hypothesized data-generating process.

Some of the most influential work on civil-military relations emphasizes the importance of structural determinants of military disloyalty (Zimmermann, 1983). When the regime lacks legitimacy, the economy is bad, or the culture is permissive of military interference in politics, the likelihood that the military will seize power is said to increase (Finer, 1988; Londregan and Poole, 1990). Additionally, Powell (2012, 1030) also argues that institutional coup-proofing measures — such as having parallel military or paramilitary forces, or an extensive security and domestic surveillance apparatus — would reduce the likelihood of coup success and finds not only that they do but that they also reduce the risk of coups.

One parameter in our model can represent some of these factors: the costs of executing the coup that the general pays. Although these costs do not affect the probability of prevailing if the coup is executed, they do affect the incentive to launch a coup because they determine the overall expected benefit of doing so. The effect of this parameter is indirect and transmitted through the way it affects the optimal resource allocation. The complete-information allocation is non-decreasing in these costs: if $m_i^*(T)$ is at the unconstrained optimum, then it is independent of the costs, and if it is at the constrained optimum, it is strictly increasing because

$S_i^*(T)$ is. In other words, when the structural factors increase the costs of a coup, the ruler can safely provide the military with larger allocations, which are helpful against the external threat.

It can be demonstrated that this means that both the safe and the risky payoffs are non-decreasing and that q^* is non-increasing. This means that under asymmetric information an increase in the costs of a coup results in generally larger allocations for the military and an expansion of the range where the ruler opts for the safe strategy (so a decrease in the ex ante risk of a coup). Thus, the expectations derived from the model are consistent with the empirical findings about the importance of structural variables.

5.4 Selecting the Generals, Part 1: Privilege

The model provides a rationale for leaders who select from privileged social groups when filling key military positions, corroborating the work of scholars who emphasize the importance of social, economic, or political-ties as determinants of military recruitment and promotion. Members of groups receiving relatively lucrative benefits from the status-quo political arrangement have less incentive to overthrow the regime, making them attractive candidates for positions within the military. It can be shown that increasing these benefits has essentially the same effect as increasing the costs of a coup. Thus, increasing \bar{b} allows the ruler to provide more resources to the military without increasing the risk of a coup. Under asymmetric information, this results in non-decreasing payoffs from both the safe and the risky strategies, as well as non-increasing q^* (that is, an expanding region where the ruler opts for the safe strategy).

The model leads us to expect that rulers who engage in exclusionary selection practices and choose their military from a restricted group of privileged elites will face a lower risk of a coup and will do better against external threats. Conversely, rulers who for some reason are unable to limit their selection to such a group but must admit representatives of other, less privileged, groups would have a higher risk of a coup and will generally perform worse against external threats.

It is not difficult to find examples of rulers engaging in exclusionary practices when it comes to their militaries. After the 1965 coup and 1966 dissolution of the monarchy in Burundi, most of the 17 military officers in the National Revolutionary Council came from the ruling Tutsi minority, while only three belonged to the Hutu ethnic majority (Kaufman and Haklai, 2008, 752). In South Africa during Apartheid, whites dominated the military leadership and were the only group allowed to fulfill combat roles. Non-whites serving in the military were relegated to supporting positions, such as making food or fixing equipment, that did not provide direct access to coercive force (Enloe, 1975, 24).

It is perhaps more interesting to compare the model's expectations to an empirical study that uses a different explanatory mechanism to derive its hypotheses. Roessler

(2011) argues that coups in sub-Saharan Africa are triggered by an internal security dilemma that arises out of the inability of elites to commit to cooperating with each other to maintain their hold on power. Rulers suspicious of the loyalty of some elites take precautionary coup-proofing measures that increase the anxiety of these elites, which makes them more prone to violence. If the ruler succeeds in excluding these elites from the coercive apparatus, this violence takes the form of a civil war, but if the ruler fails to exclude them, the violence takes the form of a coup. Because rulers cannot observe loyalty directly, they use ethnicity as an informational shortcut, “an expedient mechanism to eradicate perceived enemies at a time of high uncertainty” (Roessler, 2011, 313).

Our model has no concept of loyalty as an attribute of the potential general. Instead, loyalty is represented by the decision not to execute a coup, making it an endogenous quantity that is determined by the incentives generated from the combination of resources, competence, benefits, and threat environment. Thus, the model does not allow for the use of ethnicity as a cue for loyalty. If, however, ethnicity is a proxy for privilege (as it would be in most cases), the model rationalizes exclusionary ethnic practices simply as a way to ensure higher status quo benefits for the general.⁶ In this way, the model yields predictions that are remarkably similar to the internal security dilemma story: rulers will attempt to select the commanders of their coercive apparatus mostly from their own ethnic group (which would be privileged in other ways), and in those cases they will be at a lower risk of a coup.

One might also want to think of some additional insights provided by our mechanism. For instance, Roessler (2011, 314–16) makes a compelling argument that while ethnic exclusion can “terminate the internal security dilemma. . . it leaves the regime vulnerable to a future civil war.” One might wonder why rulers would make such a trade-off, especially because Roessler (2011, 314–15) simply asserts that civil war somehow “poses less of a threat to their political supremacy.” In contrast, our model suggests straightforward reasons for such a substitution effect. If elites are known to be disgruntled following their exclusion, they will represent a larger known threat to the regime. As we have seen, under these circumstances, the ruler will respond with increased military spending while simultaneously facing a lower risk of a coup from within. Because the larger threat has a “circling the wagons” effect that lowers the incentives for a coup, the ruler can counter the external threat more effectively. A simple extension of our model that allows the ruler to exacerbate that threat shows that the ruler does have a very strong incentive to do so (McMahon, 2014). In other words, our mechanism can explain a crucial trade-off that is not part of the internal security dilemma mechanism and as such needs to be asserted to make that explanation work.

⁶See Esteban and Ray (2008) for reasons conflict might arise along ethnic rather than class lines.

5.5 Selecting the Generals, Part 2: Competence

Many studies suggest that agents chosen on the basis of their ties to the ruler are less effective, since perceived loyalty is emphasized over merit and competence during the selection process (Quinlivan, 1999; Brooks, 2008; Gaub, 2013). As we noted above, we regard loyalty as a consequence rather than an attribute, although one might wish to consider the use of a privileged group as a measure of ties to the ruler. The model very clearly shows that the commonly argued trade-off between loyalty and competence imposes a false choice on rulers, whose optimal strategy is to select the most competent general while simultaneously increasing the probability that he will remain loyal.⁷ Moreover, it is precisely the agent's ties to the regime that permit rulers to endow the military with additional resources. One should not be surprised to learn that when the Syrian military received advanced T-72 tanks, these weapons systems were distributed first to units deemed to be closest to the Assad regime because they were led by co-ethnic Alawites and sometimes even by members of the Asad family (Bennett, 2001; Quinlivan, 1999, 147). Since more competent agents that command more resources are better positioned to deal with external threats, the model also contradicts the notion that these militaries must be of low quality.⁸

Since a key hypothesis that emerges from our analysis is that rulers always want to hire the most competent generals, we should like to take a closer look at a famous instance that seems to contradict that claim: Saddam Hussein's choice of high-ranking military officers. As we noted above, the Iraqi president exerted control over the appointments of his military commanders, giving priority to groups with close ties to the regime. In particular, Hussein favored those with whom he shared common traits — mainly fellow Ba'athists and Sunni-Arabs, as well as privileged men from the area around his hometown of Tikrit — when choosing personnel for particularly sensitive tasks (al-Marashi, 2002).

When Iraq invaded Iran in September 1980, many of these loyalist officers proved to be incompetent military leaders, resulting in a painstakingly slow advance into Iranian territory. The sluggish pace of the advance allowed the Iranian military, still reeling in the aftermath of the 1979 Islamic Revolution, the time to coalesce into a force capable of pushing the Iraqis backwards (Pelletiere and Johnson, 1991; Hiro, 1991). Hussein reacted by replacing many of these commanders, and as a result the performance of the Iraqi armed force markedly improved.

It would be easy to use this case as an example of a ruler privileging loyalty over competence when selecting military agents, and the deleterious effects of this type

⁷This is true as long as the privileged group also contains competent agents, a likely scenario given that privilege often results in access to better education and healthcare.

⁸Egorov and Sonin (2011) argue that rulers purposefully select incompetent agents as a way to minimize their exposure to the risk of a coup. Our model considers the potential risks associated with incompetence, and draws much different conclusions.

of decision-making calculus. However, it is important to dig a bit deeper to understand the decision-making calculus of the Hussein regime, and the consequences of these decisions for military effectiveness. First of all, evidence suggests that Saddam Hussein was selecting for both competence and privilege when appointing officers to military positions prior the war. Woods et al. (2011, 14) write that “At the war’s outset, Saddam was heavily influenced by Ba’ath ideology. He believed that any Ba’ath leader could, at the same time, be a competent military commander.” In fact, the regime had a famous slogan: “al-askari al-jayyid huwa al-Baathi al-jayyid”, which means “the good military man is the good Baathist” (Parasiliti and Antoon, 2000, 134).

In other words, not only did Hussein not regard the selection from the privileged Baathist elite as some sort of substitute for competence, he seems to have thought that membership in the party was a good indicator of high military competence. His behavior is thus in line with the model’s expectations.

As it turned out, Hussein was mistaken about the direct relationship between membership in the privileged group and competence. Since our model assumes that competence is directly observable, it does not allow for such mistakes. In its present form, the model cannot account for Hussein’s initial choice. However, if our model is right, then Hussein’s intention must have been to select competent generals. This implies that upon realizing that he had made a mistake, Hussein should immediately have moved to correct it by making appropriate replacements. Since performance in war can be regarded as a direct test of competence, the fact that Hussein did replace unsuccessful commanders with successful ones can be taken as evidence that supports our model. Moreover, the fact that Hussein continued to select from the privileged group lends further support to our model and undermines the idea of a trade-off between loyalty and competence.⁹

We can take our analysis further and turn Hussein’s apparent mistake to our analytical advantage. Some studies suggest that while the loyalty-competence trade-off is real, large external threats can swamp the fear of a coup and cause the ruler to focus more on competence rather than loyalty (Talmadge, 2013). If such a decision is conditional on the level of threat, then the choice of competent commanders must be transient: The diminution of the threat must cause the ruler to revert to form. In this instance, after the war’s end Hussein would be expected to replace the competent, but now dangerous, generals with incompetent cronies. Our model, on the other hand, would lead us to expect precisely the opposite outcome because the incentive to select the most competent general is independent of the level of threat. Since performance in the war has allowed him to identify the competent commanders, Hussein would be expected to retain them after the war’s end.

In order to assess these divergent expectations, we identify the senior leaders of

⁹Pelletiere and Johnson (1991, 59) note this trend: “Most of Iraq’s higher level commanders appear to have been politically reliable professionals after 1982. Indeed, from 1984 on, the issue of competence seems to have been the principal deciding factor for advancement.”

the Iraqi armed forces — members of the high command — during the last two years of the Iran-Iraq War (1987 and 1988). The high command is inclusive to a variety of senior military leadership positions, from commander-in-chief Saddam Hussein and the minister of defense to the commanders of the navy and the seven army corps (Bengio, 1989, 1990). After determining the individuals who held these key posts, we tracked their career trajectories for the first few years after the war to determine if they were purged.

In all, we were able to track the post-war career trajectories for 23 of the 27 members of the 1987/88 military high command (see online Appendix B). We excluded Hussein himself (the 28th commander) from this analysis. The data reveal that Hussein continued to employ a sizable majority of his senior military leaders after the war. As is shown in Table 1, almost three-quarters of the generals continued to hold the same position or were promoted by the regime. This evidence, while it is only suggestive, provides support for the predictions of our theory: Hussein kept his war-proven military commanders after the war. There is no post-war information for four of the generals who served in the high command, possibly because they were purged by the regime. Yet even if we were to assume that all four were purged, it would mean that Hussein retained 76 percent of the still-living members of the high command from 1987/88 for at least two years after the war.

Outcome	Number	Relative Freq.
Held Same Position	11	0.41
Were Promoted	8	0.30
Death in Combat	1	0.04
Death in Accident	1	0.04
Retirement	2	0.07
Unknown	4	0.15
Total	27	1.01

Notes: Career trajectories for the men who held positions in the Iraqi military high command during 1987 or 1988. Outcomes determined by the status of each individual between the end of the Iran-Iraq war and 1990/1.

Table 1: 1987-88 Iraqi Military High Command After Iran-Iraq War

Hussein’s actions clearly demonstrate a ruler who was seeking both to maximize competence and ensure loyalty while selecting military officers. The Iraqi regime used the crucible of war to identify competent commanders, and then continued to rely on these individuals after the fighting had ended. While our evidence is limited to the senior leadership of the Iraqi armed forces, we believe that the calculus driving the decisions of rulers like Hussein is most evident among these individuals, who are both highly influential and visible within the state. Focusing on the high

command also allows us to avoid confusing political calculations with the downsizing that occurs among the ranks of most militaries following long and costly wars.

6 Conclusion

Examining the Guardianship Dilemma allows us to make several contributions to the study of political instability in general, and civil-military relations more specifically. First, the same threats that necessitate the creation of a military for defense also help to keep this force loyal. Ultimately, the pernicious effects of the Guardianship Dilemma are due primarily to rulers' uncertainty about the threat environment, rather than the severity of these threats. Second, rulers in coup-prone states are better off by staffing their militaries with persons from privileged groups, since regimes can actually *increase* the fighting power of their armed forces through selective appointments. If individuals with a higher-status-quo benefit have less incentive to overthrow the regime, they can be trusted with more coercive power. Third, rulers also select generals on the basis of their competence, since a competent commander can more effectively manage the state's military resources. While this competence increases the likelihood that a coup will succeed if attempted, rulers prefer to manage the generals' loyalty by restricting the resources that flow to militaries, rather than by appointing dolts who will waste what they are given.

While the model enables us to characterize the most basic features of the dilemma, further research is necessary to extend and refine these insights. A useful first step would be to examine whether the Guardianship Dilemma depends on the source of the threat facing the state, especially in the context of foreign versus domestic enemies. While previous research predicts that civil-military relations differ drastically on the basis of whether the threat comes from within or outside of states' borders (Desch, 1999), the dynamics we outline exist in the presence of any threat that endangers the government. Dissimilar threats may, however, differentially affect the variables that impinge upon the dilemma, changing the extent to which rulers must relax or impose constraints on their armed forces. There may also exist levels of uncertainty about the strategic environment that vary systematically across unlike threats. While these differences are likely to be manifest in the empirical record, it would be worthwhile to explore whether or not this variation in behavior is due to same underlying mechanism.

Another step would be to analyze an extension to the model in which the interaction between military and political elites takes place over multiple periods. The game in our model ends when the external threat is faced. In practice, of course, whoever is in charge of the state must continue to rule. If the military has been endowed with the resources to face a threat, the sudden termination of this loyalty-inducing threat could put the regime in grave danger, leading rulers to prefer endur-

ing hostilities with opponents. In this, it may be possible to develop a civil-military logic for the cultivation of rivalries. One could also explore how the actions of rulers might change if the competence of potential generals is hard to observe, as it was initially for Saddam Hussein. The model reveals that knowing the competence of the military leadership is important for rulers who are calibrating an appropriate response to external threats and the risk of a coup, and that, consequently, uncertainty about competence could potentially be costly or dangerous (or both!) for the regime. As a result, we may identify rulers who invest in military academies, conduct exercises, or use the process of fighting in conflict to separate the types of military officers in their armed forces.

References

- Acemoglu, Daron, Davide Ticchi and Andrea Vindigni. 2008. "A Theory of Military Dictatorships." *NBER Working Paper Series, No. 13915* .
- al-Marashi, Ibrahim. 2002. "Iraq's Security and Intelligence Network: A Guide and Analysis." *Middle East Review of International Affairs (MERIA) Journal* 6(3):1–13.
- Arbatli, Cemal Eren and Ekim Arbatli. 2014. "External threats and political survival: Can dispute involvement deter coup attempts?" *Conflict Management and Peace Science* .
- Bengio, Ofra. 1989. Iraq. In *Middle East Contemporary Survey, 1987*, ed. Itamar Rabinovich, Haim Shaked and Ami Ayalon. Vol. XI Boulder, CO: Westview Press pp. 423–59.
- Bengio, Ofra. 1990. Iraq. In *Middle East Contemporary Survey, 1988*, ed. Ami Ayalon and Haim Shaked. Vol. XII Boulder, CO: Westview Press pp. 500–43.
- Bennett, Richard M. 2001. "The Syrian Military: A Primer." *Middle East Intelligence Bulletin* 3(8).
- Besley, Timothy and James A Robinson. 2010. "Quis Custodiet Ipsos Custodes? Civilian Control Over the Military." *Journal of the European Economic Association* 8(2-3):655–63.
- Brooks, Risa A. 1998. Political-Military Relations and the Stability of Arab Regimes. In *Adelphi Paper 324*. New York, NY: Oxford University Press.
- Brooks, Risa A. 2008. *Shaping Strategy: The Civil-Military Politics of Strategic Assessment*. Princeton, NJ: Princeton University Press.
- Desch, Michael C. 1999. *Civilian Control of the Military: The Changing Security Environment*. Baltimore, MD: Johns Hopkins University Press.
- Egorov, Georgy and Konstantin Sonin. 2011. "Dictators and their Viziers: Endogenizing the Loyalty-Competence Trade-Off." *Journal of the European Economic Association* 9(5):903–30.
- Enloe, Cynthia H. 1975. "Ethnic Factors in the Evolution of the South African Military." *Issue: A Journal of Opinion* 5(4):21–28.
- Esteban, Joan and Debraj Ray. 2008. "On the Saliency of Ethnic Conflict." *American Economic Review* 98(5):2185–2202.

- Feaver, Peter. 1996. "The Civil-Military Problematique: Huntington, Janowitz, and the Question of Civilian Control." *Armed Forces and Society* 23(2):149–78.
- Feaver, Peter. 1999. "Civil-Military Relations." *Annual Review of Political Science* 2:211–41.
- Finer, S. E. 1988. *The Man on Horseback*. Boulder, CO: Westview Press.
- Gaub, Florence. 2013. "The Libyan Armed Forces between Coup-proofing and Repression." *Journal of Strategic Studies* 36(2):221–44.
- Goemans, Henk E., Kristian Skrede Gleditsch and Giacomo Chiozza. 2009. "Introducing Archigos: A Dataset of Political Leaders." *Journal of Peace Research* 46(2):269–83.
- Hiro, Dilip. 1991. *The Longest War*. New York, NY: Routledge.
- Horowitz, Donald. 1985. *Ethnic Groups in Conflict*. Berkeley, CA: University of California Press.
- Huntington, Samuel P. 1957. *The Soldier and the State: The Theory and Practice of Civil-Military Relations*. Boston, MA: Belknap.
- Kaufman, Eric and Oded Haklai. 2008. "Dominant ethnicity: from minority to majority." *Nations and Nationalism* 14(4):743–767.
- Londregan, John and Keith Poole. 1990. "Poverty, the coup trap and the seizure of executive power." *World Politics* 42(2):151–83.
- Lutterbeck, Derek. 2013. "Arab Uprisings, Armed Forces, and Civil-Military Relations." *Armed Forces and Society* 39(1):28–52.
- Luttwak, Edward. 1979. *Coup d'état, a practical handbook*. Cambridge, MA: Harvard University Press.
- McMahon, R. Blake. 2014. "Circling the Wagons: Civil-Military Relations and International Disputes." Presented at the Meeting of the Peace Science Society, Philadelphia.
- Neff, Donald. 1994. "Israel-Syria: Conflict at the Jordan River, 1949-1967." *Journal of Palestine Studies* 23(4):26–40.
- Nossiter, Adam. 2012. "Leaders of Mali's Military Coup Seem to Have Uncertain Grasp on Power." *The New York Times*, March 24, A8.
- Parasiliti, Andrew and Sinan Antoon. 2000. "Friends in Need, Foes to Heed: The Iraqi Military in Politics." *Middle East Policy* 7(4):130–40.

- Pelletiere, Stephen C. and Douglas V. Johnson. 1991. *Lessons Learned: The Iran-Iraq War*. Carlisle Barracks, PA: Strategic Studies Institute, U.S. Army War College.
- Pilster, Ulrich and Tobias Boehmelt. 2011. "Coups-Proofing and Military Effectiveness in Interstate Wars, 1965-1999." *Conflict Management and Peace Science* 28(4):331–50.
- Pipes, Daniel. 1989. "The Alawi Capture of Power in Syria." *Middle Eastern Studies* 25(4):429–50.
- Powell, Jonathan. 2012. "Determinants of the Attempting and Outcome of Coups d'Etat." *Journal of Conflict Resolution* 56(6):1017–40.
- Powell, Jonathan M. and Clayton L. Thyne. 2010. "Global Instances of Coups from 1950 to 2009: A New Dataset." *Journal of Peace Research* 48(2):249–59.
- Quinlivan, James T. 1999. "Coups-Proofing: Its Practice and Consequences in the Middle East." *International Security* 24(2):131–65.
- Roessler, Philip. 2011. "The Enemy Within: Personal Rule, Coups, and Civil War in Africa." *World Politics* 63(2):300–46.
- Slantchev, Branislav. 2003. "The Principle of Convergence in Wartime Negotiations." *American Political Science Review* 97(4):621–32.
- Staniland, Paul. 2008. "Explaining Civil-Military Relations in Complex Political Environments: India and Pakistan in Comparative Perspective." *Security Studies* 17(2):322–62.
- Sutter, Daniel. 2000. "A Game-Theoretic Model of the Coup D'Etat." *Economics & Politics* 12(2):205–23.
- Svolik, Milan W. 2009. "Power Sharing and Leadership Dynamics in Authoritarian Regimes." *American Journal of Political Science* 53(2):477–94.
- Svolik, Milan W. 2012. *The Politics of Authoritarian Rule*. New York, NY: Cambridge University Press.
- Svolik, Milan W. 2013. "Contracting on Violence: The Moral Hazard in Authoritarian Repression and Military Intervention in Politics." *Journal of Conflict Resolution* 57(5):765–94.
- Talmadge, Caitlin. 2013. "The Puzzle of Personalist Performance: Iraqi Battlefield Effectiveness in the Iran-Iraq War." *Security Studies* 22(2):180–221.

- Talmadge, Caitlin and Varun Piplani. 2014. "Does War Help Civil Military Relations? Prolonged Interstate Conflict and the Risk of Coup Attempts." Working Paper.
- Thompson, Leonard. 2001. *A History of South Africa*. New Haven, CT: Yale University Press.
- Woods, Kevin M., Williamson Murray, Elizabeth A. Nathan, Laila Sabara and Ana M. Venegas. 2011. *Saddam's Generals: Perspectives of the Iran-Iraq War*. Alexandria, VA: Institute for Defense Analysis.
- Zimmermann, Ekkart. 1983. *Political Violence, Crises, and Revolutions*. Boston, MA: G.K. Hall and Co.

TO BE MADE AVAILABLE ONLINE

Appendix A: Proofs

Proof of Lemma 1. If G remains loyal and does not mount a coup, this threat is realized and G 's expected payoff is $p(m, T; \theta)b_i$. If a coup did occur but it failed, G is eliminated entirely so his payoff once the threat is realized remains 0. If a coup succeeded, G obtains the benefit of rule and fights the external threat (R 's security resources are assumed lost and unavailable to G). Thus, G 's expected payoff from a coup is $p(m, 1; \theta)p(m, T; \theta) - c$. By subgame-perfection, G will remain loyal if

$$p(m, T; \theta)[p(m, 1; \theta) - b_i] < c,$$

execute a coup if the strict inequality is reversed, and be indifferent otherwise. We can rewrite this as $T < T_i^*(m, \theta)$, where the latter is defined in (1). This establishes the sufficiency part of the claim.

Letting $x \equiv \theta m$, we can observe that

$$\frac{dT_i^*}{dx} = \left(\frac{1}{c}\right) \left[1 - (b_i + c) - \frac{1}{(1+x)^2}\right],$$

which means that

$$\text{sgn}\left(\frac{dT_i^*}{dx}\right) = \text{sgn}\left(1 - (b_i + c) - \frac{1}{(1+x)^2}\right).$$

This yields a quadratic, $x^2 + 2x - \frac{b_i+c}{1-(b_i+c)} > 0$, which is a parabola that opens up. Although the discriminant is $4/(1 - (b_i + c)) > 0$, the smaller root is negative, which means that the inequality is satisfied for all

$$x > \frac{1}{\sqrt{1 - (b_i + c)}} - 1. \quad (4)$$

But now $T_i^*(m, \theta) \geq 0$ implies that

$$x \geq \frac{b_i + c}{1 - (b_i + c)} > \frac{1}{\sqrt{1 - (b_i + c)}} - 1,$$

where the second inequality is readily verified under Assumption 1, and so (4) must be satisfied whenever T_i^* is non-negative. In other words, when T_i^* is non-negative it must be increasing in both θ and m , as claimed.

To prove necessity, we need to show that there is no equilibrium where G executes a coup with positive probability when indifferent. Suppose, to the contrary, that he does execute a coup with positive probability, perhaps even certainty, when

indifferent. First, note that if R 's expected payoff from a coup is at least as good as the expected payoff from loyalty, then the fact that R strictly prefers not hiring a general to a coup also implies that R would not hire a general in this case. In other words, whenever G gets hired in equilibrium, it must be that R strictly prefers him to remain loyal:

$$p(m, T; \theta) > p(1, \theta m; 1)p(m, T; 1). \quad (5)$$

Second, we show that R can do strictly better by ensuring G 's loyalty. Letting $q \in (0, 1]$ denote the probability of a coup, R 's expected payoff is $qp(1, \theta m; 1)p(m, T; 1) + (1 - q)p(m, T; \theta) - m$. Since G is indifferent, it must be that $T_i^*(m, \theta) = T > 0$, which further implies that T_i^* is increasing in m . This now means that any $\hat{m} < m$ would result in $T_i^*(\hat{m}, \theta) < T$, ensuring G 's loyalty. Consider now some such $\hat{m} < m$ that is arbitrarily close to m , and observe that this means that $p(\hat{m}, T; \theta)$ is arbitrarily close to $p(m, T; \theta)$. By (5), $qp(1, \theta m; 1)p(m, T; 1) + (1 - q)p(m, T; \theta) < p(m, T; \theta)$ for any $q \in (0, 1]$, which means that we can always find \hat{m} such that $qp(1, \theta m; 1)p(m, T; 1) + (1 - q)p(m, T; \theta) < p(\hat{m}, T; \theta)$. In other words, R strictly prefers to reduce m by an arbitrarily small amount and ensure G 's loyalty. But this contradicts the equilibrium requirement that R 's strategy be optimal. Therefore, there can be no equilibrium where G executes a coup with positive probability when indifferent. This establishes the necessity part of the claim. ■

Proof of Lemma 2. If R 's choices avoid a coup, her payoff is $U = p(m, T; \theta) - m$, and the loyalty constraint, $T \geq T_i^*(m, \theta)$, must obtain. Solving for the constraint yields the quadratic $(1 - (b_i + c))\theta^2 m^2 - (b_i + c + cT)\theta m - cT \leq 0$, whose discriminant is $(b_i + c - cT)^2 + 4cT > 0$. Under Assumption 1 the smaller root is negative, so let $S_i^*(T)$ be the larger root defined in (2). Since the coefficient on the squared term is positive, the constraint is satisfied for all $\theta m \leq S_i^*(T)$.

R 's payoff is strictly increasing in θ and concave in m :

$$\frac{dU}{d\theta} = \frac{mT}{(\theta m + T)^2} > 0 \quad \text{and} \quad \frac{dU}{dm} = \frac{\theta T}{(\theta m + T)^2} - 1.$$

Let the solution to the first-order condition on m be defined as

$$\tilde{m}(\theta) = \max \left(0, \sqrt{\frac{T}{\theta}} - \frac{T}{\theta} \right),$$

so clearly the unconstrained maximum is at $(\tilde{m}(\bar{\theta}), \bar{\theta})$. Let $S(T) = \bar{\theta}\tilde{m}(\bar{\theta})$ be the loyalty induced if R were to provide G of maximal competence with the level of resources optimal for dealing with the threat. If this level of disloyalty does not exceed the maximum level that avoids a coup, $S(T) \leq S_i^*(T)$, then the unconstrained maximum is the unique solution to R 's maximization problem.

If $S(T) > S_i^*(T)$, then the induced level of disloyalty exceeds the safe maximum, and G would execute a coup if he were provided with such resources. Since

this cannot happen in an equilibrium where coups are avoided, the loyalty constraint must bind: $\theta m = S_i^*(T)$. (If it were slack at some θm , then R could strictly increase her payoff by increasing θ until it binds.) This means that R 's expected payoff can be written as

$$U = \frac{S_i^*(T)}{S_i^*(T) + T} - \frac{S_i^*(T)}{\theta},$$

which is strictly increasing in θ . Therefore, R will pick $\bar{\theta}$ again except that this time she will handicap G by providing him with fewer resources. ■

Proof of Lemma 3. When G 's resources are not constrained by loyalty considerations, the envelope theorem tells us that

$$\frac{dU(m_i^*(\bar{\theta}), \bar{\theta})}{d\bar{\theta}} = \frac{\partial U(m_i^*(\bar{\theta}), \bar{\theta})}{\partial \bar{\theta}} = \frac{m_i^*(\bar{\theta})T}{(\bar{\theta}m_i^*(\bar{\theta}) + T)^2} > 0.$$

Since R 's payoff when not hiring a general can be represented by the payoff of hiring a general with competence $\theta = 1$ for whom the constraint is not binding, we conclude that if $\bar{\theta} < 1$, then R strictly prefers not to hire a general than to hire one whose loyalty will not be a problem at the optimal level of resource provision. Since R 's payoff is strictly smaller when the loyalty constraint binds, this further implies that R will not want to hire a general at all. This establishes case (i) of the lemma.

If $\bar{\theta} > 1$, then R strictly prefers to hire G provided that his loyalty will not be a problem. We know, however, that for $\bar{\theta}$ sufficiently high, $S(T) > S_i^*(T)$ will obtain, and so R will be forced to reduce the resources in order to ensure G 's loyalty. Would she still wish to hire this general? Assume that $S(T) > S_i^*(T)$ so $m_i^*(T) = S_i^*(T)/\bar{\theta}$. Hiring a general yields

$$\frac{S_i^*(T)}{S_i^*(T) + T} - \frac{S_i^*(T)}{\bar{\theta}} > 0,$$

where we can establish the inequality as follows. The inequality holds if, and only if, $\bar{\theta} > S_i^*(T) + T$. But since $S(T) > S_i^*(T)$ here, it follows that $\sqrt{\bar{\theta}T} > S_i^*(T) + T$, which reduces to $\bar{\theta} > (S_i^*(T))^2/T + 2S_i^*(T) + T > S_i^*(T) + T$. Thus, whenever the loyalty constraint binds, R 's (constrained) payoff is strictly positive.

Not hiring a general with optimal allocation $m = \sqrt{T} - T$ (provided $T < 1$) yields

$$\frac{\sqrt{T} - T}{\sqrt{T}} - \sqrt{T} + T = 1 + T - 2\sqrt{T} > 0.$$

Since $T \geq 1$ means that not hiring yields a payoff of zero (because the optimal allocation is zero), it follows that in all such cases R strictly prefers to hire a general even if doing so requires R to impose constraints on him. This establishes case (ii) of the lemma.

Suppose then that $T < 1$, so that the payoffs from hiring and not hiring are both positive. We now show that it is possible that R prefers not to hire at all. Note first that

$$\lim_{c \rightarrow 0} S_i^*(T) = \frac{b_i}{1 - b_i},$$

and since we require that $S_i^*(T) < S(T)$, the condition that the constraint is binding will be satisfied for any

$$b_i < \frac{S(T)}{1 + S(T)}.$$

This means that as $b_i \rightarrow 0$, the constraint must be binding, and since $\lim_{b_i \rightarrow 0} b_i/(1 - b_i) = 0$, we obtain

$$\lim_{c \rightarrow 0, b_i \rightarrow 0} \frac{S_i^*(T)}{S_i^*(T) + T} - \frac{S_i^*(T)}{\bar{\theta}} = 0 < 1 + T - 2\sqrt{T}.$$

In other words, if c and b_i are sufficiently small, then it must be the case that R strictly prefers not to hire. This establishes case (iii) of the lemma. ■

Proof of Lemma 4. It is clear by inspection of (2) that S_i^* is strictly increasing in b_i . Since $S(T)$ is constant in b_i , it follows that $b^* > 0$ such that $S_i^*(T) = S(T)$ exists and is unique. If $\bar{b} \leq b^*$, then the loyalty constraint is binding, so the military allocation is $m_i^*(T) = S_i^*(T)/\bar{\theta}$, which is increasing in $S_i^*(T)$. Moreover, since this constrained allocation is less than the unconstrained optimum, it follows that R 's expected payoff is strictly increasing in m_i^* as well. In other words, in this case R 's expected payoff strictly increases in b_i , which implies that she must pick \bar{b} . If $\bar{b} > b^*$, then the loyalty constraint is no longer binding, so R 's military allocation is at the unconstrained optimum, which itself is independent of b_i . In these cases, R is indifferent among any $b_i \in (b^*, \bar{b}]$, as claimed. ■

Proof of Lemma 5. Suppose a coup will occur, so R 's payoff is

$$U = \left(\frac{m}{1 + \theta m} \right) \left(\frac{q}{m + T_S} + \frac{1 - q}{m + T_L} \right) - m,$$

which is always strictly worse than not hiring a general for any $m > 0$. Since $T_i^*(0, \theta) = 0 < T_S$, the probability of a coup is zero when $m = 0$, which implies that in any subgame where a coup is certain to occur it must be the case that $m > 0$, and so R is strictly better off not hiring a general. In other words, there exists no equilibrium where a coup is certain to occur. ■

Proof of Lemma 6. NO COUP. Suppose there is an equilibrium in which no coups occur regardless of the size of the threat. We know that this requires m to be such that G remains loyal under T_S . It turns out that $m_i^*(T_S)$ must be the optimal security-preserving allocation under asymmetric information as well. We know that it cannot exceed that level because if it did, G would execute a coup under T_S . It also cannot be less than that level because if it did, R 's payoffs under both T_S and T_L (under Assumption 3) would decrease, leading to a decrease in the expected payoff as well. Thus, the best expected payoff that R can obtain where no coup occurs is

$$U_N(q) = q \left(\frac{\theta m_i^*(T_S)}{\theta m_i^*(T_S) + T_S} \right) + (1 - q) \left(\frac{\theta m_i^*(T_S)}{\theta m_i^*(T_S) + T_L} \right) - m_i^*(T_S).$$

Since $m_i^*(T_S)$ does not depend on q , U_N is a simple linear function of q . In particular, since $T_S < T_L$, it is strictly increasing

$$\frac{dU_N}{dq} = \frac{\theta m_i^*(T_S)(T_L - T_S)}{(\theta m_i^*(T_S) + T_S)(\theta m_i^*(T_S) + T_L)} > 0.$$

We now show that if $\theta \leq 1$, then R prefers to go it alone when the alternative is hiring a general who would not execute a coup. This follows immediately from the fact that $\theta < 1 \Rightarrow U_A > U_N$ for any $m > 0$ and any q . We can write $U_A > U_N$ as

$$q [p(m, T_S; 1) - p(m, T_S; \theta)] + (1 - q) [p(m, T_L; 1) - p(m, T_L; \theta)] > 0,$$

so it is sufficient to show that both bracketed terms are positive. Since $p(m, T; \theta)$ is strictly increasing in θ , they are positive when $\theta < 1$, so the claim holds. Moreover, since $\theta \leq T_S$ implies that $m_i^*(T_S) = 0$, we obtain $U_N = 0 < U_A$, so R will also prefer to go it alone in this case as well. Thus, the necessary condition for hiring G in such an equilibrium is $\theta > \max(1, T_S)$.

PROBABILISTIC COUP. Suppose there is an equilibrium in which G executes a coup under T_S but remains loyal under T_L . This means that $T_S < T_i^*(m, \theta) \leq T_L$. Recalling from Lemma 1 that T_i^* is increasing in both parameters whenever it is positive (as it must be here), we conclude that the optimal allocation must be some $m_C \in (m_i^*(T_S), m_i^*(T_L)]$.

When the coup is probabilistic, R 's expected payoff is

$$U_C(q) = q \left[\left(\frac{1}{1 + \theta m} \right) \left(\frac{m}{m + T_S} \right) - m \right] + (1 - q) \left(\frac{\theta m}{\theta m + T_L} - m \right). \quad (6)$$

We now show that $\theta \leq 1 \Rightarrow U_A > U_C$, so R will never hire a general that is less competent than herself if she expects the continuation game to involve a probabilistic coup. We can write $U_A > U_C$ as

$$q [p(m, T_S; 1) - p(1, \theta m; 1)] + (1 - q) [p(m, T_L; 1) - p(m, T_L; \theta)] > 0,$$

so it is sufficient to show that both bracketed terms are positive. The first is positive because $p(1, \theta m; 1) < 1$, and the second is non-negative if $\theta \leq 1$ because $p(m, T; \theta)$ is strictly increasing in θ . Moreover, $\theta \leq T_S < T_L$ implies that $m_i^*(T_S) = m_i^*(T_L) = 0$, so there exists no m_C that will induce a probabilistic coup. In other words, if $\theta \leq T_S$, then such an equilibrium does not exist. Thus, the necessary condition for hiring G in such an equilibrium is also $\theta > \max(1, T_S)$.

Since R will not hire G with $\theta \leq 1$, for the remainder of this proof we shall assume that $\theta > 1$. The unconstrained FOC for (6) is

$$\begin{aligned} \frac{\partial U_C}{\partial m} &= \frac{q(T_S - \theta m^2)}{(1 + \theta m)^2(m + T_S)^2} + \frac{(1 - q)\theta T_L}{(\theta m + T_L)^2} - 1 \\ &= q \left[\frac{T_S - \theta m^2}{(1 + \theta m)^2(m + T_S)^2} - \frac{\theta T_L}{(\theta m + T_L)^2} \right] + \frac{\theta T_L}{(\theta m + T_L)^2} - 1 \\ &\equiv q\zeta + \frac{\theta T_L}{(\theta m + T_L)^2} - 1 = 0. \end{aligned} \quad (7)$$

Since the derivative is strictly decreasing in m , it attains a maximum at $m = 0$, where it is strictly positive if, and only if, $qT_L + (1 - q)\theta T_S > 1$. By Assumption 2 and $\theta > 1$, this condition is satisfied, so the fact that $\lim_{m \rightarrow \infty} \frac{\partial U}{\partial m} = -1$ implies that there exists a unique $m_C(q) > 0$ for which the FOC is satisfied (i.e., the function is concave). The question now is to ensure that the solution satisfies the constraints.

We begin by showing that $m_C(q)$ must be decreasing. The implicit function theorem tells us that (7) implies that

$$\frac{dm_C}{dq} = - \frac{\partial^2 U_C}{\partial m \partial q} / \frac{\partial^2 U_C}{\partial m \partial m_C}$$

which then tells us that since

$$\frac{\partial^2 U_C}{\partial m \partial m_C} < 0 \Rightarrow \text{sgn} \left(\frac{dm_C}{dq} \right) = \text{sgn} \left(\frac{\partial^2 U_C}{\partial m \partial q} \right) = \text{sgn}(\zeta) = \text{sgn} \left(1 - \frac{\theta T_L}{(\theta m + T_L)^2} \right),$$

where the last step also follows from (7) and $q > 0$. This, of course, yields

$$\text{sgn} \left(1 - \frac{\theta T_L}{(\theta m + T_L)^2} \right) = -1 \Leftrightarrow m < \sqrt{\frac{T_L}{\theta}} - \frac{T_L}{\theta} \equiv \tilde{m},$$

where the last expression is the unconstrained optimum for the complete-information case under T_L .

We now show that m_C can never exceed this value. Consider the payoff in (6). The expression in the square brackets (the expected payoff from a coup with T_S) is strictly decreasing in m because

$$\frac{T_S - \theta m^2}{(1 + \theta m)^2(m + T_S)^2} - 1 < 0$$

obtains. To see this, observe that it is certainly true for any $T_S - \theta m^2 \leq 0$. When this expression is positive, we can write the inequality as $T_S - \theta m^2 < (1 + \theta m)^2(m + T_S)^2$, and observe that the left-hand side is strictly decreasing in m while the right-hand side is strictly increasing. Thus, if the inequality holds at $m = 0$, it must hold at $m > 0$ as well. But at $m = 0$ the inequality reduces to $T_S < T_S^2 \Leftrightarrow 1 < T_S$, which holds by Assumption 2. Thus, the first component in the expected payoff is always strictly decreasing in m .

The second component of this payoff is, of course, the complete-information payoff without a coup against T_L , and we know that its unconstrained optimum is $\tilde{m} = \sqrt{T_L/\theta} - T_L/\theta$. This immediately tells us that $m_C < \tilde{m}$: if this were not so, one could improve the payoff by decreasing m_C to \tilde{m} since this will strictly increase both components.

Thus, $m_C(q) < \tilde{m}$, which in turn means that $\text{sgn}(\zeta) = -1$, and we conclude that $m_C(q)$ is strictly decreasing.

Observe now that at $q = 0$, the payoff in (6) is equivalent to the complete-information case under T_L , which means that $m_C(0) = m_i^*(T_L) > m_i^*(T_S)$, where the inequality follows from Assumption 3, so the constraints are satisfied (the general executes a coup if the threat is T_S but does not if it is T_L). Moreover, since $m_i^*(T_L)$ is the (possibly constrained) optimum against T_L , it follows that

$$U_C(0) = \frac{\theta m_i^*(T_L)}{\theta m_i^*(T_L) + T_L} - m_i^*(T_L) > \frac{\theta m_i^*(T_S)}{\theta m_i^*(T_S) + T_L} - m_i^*(T_S) = U_N(0),$$

which means that at $q = 0$, the ruler must strictly prefer to play the risky strategy by endowing G with enough resources to meet the large external threat. (Of course, at $q = 0$, this risk is zero.)

Consider now what happens as q increases, in which case we have shown that m_C must decrease. There are two cases, depending on whether $m_C(q)$ satisfies the constraints or not.

Case 1: $m_C(q) \geq m_i^*(T_L)$, which implies that the solution must be constrained at $m_i^*(T_L)$ (or else G would execute the coup regardless of the threat size): since the payoff function is concave in m , it must be increasing for all $m < m_C(q)$. Moreover, since $m_C < \tilde{m}$, it follows that $m_C(q) \geq m_i^*(T_L)$ can only obtain when $m_i^*(T_L)$ is the constrained solution to the complete-information case, which means that $m_i^*(T_L) = S_i^*(T_L)/\theta$. For U_C to be decreasing, it must be the case that

$$\begin{aligned} \frac{dU_C}{dq} &= \frac{\partial U_C}{\partial m_i^*(T_L)} \frac{dm_i^*(T_L)}{dq} + \frac{\partial U_C}{\partial q} = \frac{\partial U_C}{\partial q} \\ &= \frac{m_i^*(T_L)}{(1 + \theta m_i^*(T_L))(m_i^*(T_L) + T_S)} - \frac{\theta m_i^*(T_L)}{\theta m_i^*(T_L) + T_L} < 0, \end{aligned}$$

where the first step follows from the fact that $\frac{dm_i^*(T_L)}{dq} = 0$ at the constrained solution. Letting $m \equiv m_i^*(T_L) > 0$ to simplify notation, we can rewrite the inequality

above as

$$\frac{1}{(1 + \theta m)(m + T_S)} < \frac{\theta}{\theta m + T_L}, \quad (8)$$

Recall that $m_i^*(T_L)$ is the constrained solution to the complete information case, which means that $S(T_L) > S_i^*(T_L) > 0$, which in turn implies that $S(T_L) > 0$ must be satisfied, and so $\theta > T_L$ must obtain. But this now implies that

$$\frac{1}{(1 + \theta m)(m + T_S)} < \frac{1}{(1 + mT_L)(m + T_S)} \quad \text{and} \quad \frac{\theta}{\theta m + T_L} > \frac{T_L}{mT_L + T_L} = \frac{1}{1 + m},$$

so it will be sufficient to show that

$$\frac{1}{(1 + mT_L)(m + T_S)} < \frac{1}{1 + m} \quad \Leftrightarrow \quad 1 + m < (1 + mT_L)(m + T_S),$$

where the last inequality is easily verified because $mT_L > 0$ and $T_S > 1$ together imply that $(1 + mT_L)(m + T_S) > m + T_S > m + 1$. Thus, U_C is strictly decreasing in q for any $m_C \geq m_i^*(T_L)$.

Summarizing, start with $q = 0$, where the solution is $m_C = m_i^*(T_L)$. If $m_i^*(T_L)$ is the constrained solution to the complete-information case, then it is possible that the solution to (7) is actually strictly greater. If this is so, then increasing q will decrease this solution until at some point it will equal $m_i^*(T_L)$: in this interval, the optimal allocation is constant at $m_i^*(T_L)$, and the payoff is strictly decreasing. If $m_i^*(T_L)$ is the unconstrained solution, then the fact that $m_C(q)$ is decreasing means that the second case applies.

Case 2: $m_C(q) \in [m_i^*(T_S), m_i^*(T_L)]$. In this region, the constraint that ensures that G remains loyal under T_L is no longer binding, and since this means that $\frac{\partial U_C}{\partial m} = 0$ at the optimum, we can apply the envelope theorem to obtain

$$\begin{aligned} \frac{dU_C}{dq} &= \frac{\partial U_C}{\partial m} \frac{dm}{dq} + \frac{\partial U_C}{\partial q} = \frac{\partial U_C}{\partial q} \\ &= \frac{m_C}{(1 + \theta m_C)(m_C + T_S)} - \frac{\theta m_C}{\theta m_C + T_L} < 0, \end{aligned}$$

where we can establish this inequality as follows. If $m_i^*(T_L)$ is the constrained solution to the complete-information case, then $\theta > T_L$ must obtain, and the argument following (8) applies. If, on the other hand, $m_i^*(T_L)$ is the unconstrained solution to the complete-information case, then we argue as follows. Loosely, since the first component of the payoff in (6) is strictly decreasing in m while the second one is strictly increasing, putting more weight on the first component decreases m_C (we showed this already), which in turn decreases U_C . We need to show that

$$\frac{m_C}{(1 + \theta m_C)(m_C + T_S)} - m < \frac{\theta m_C}{\theta m_C + T_L} - m.$$

Recall that the left-hand side is strictly decreasing in m and we know that the right-hand side is strictly increasing because m_C is smaller than the unconstrained optimum of the complete-information case under T_L . But since at $m = 0$ both sides are zero, the inequality must obtain for any $m > 0$ in this region. In other words, U_C is strictly decreasing here as well. Note in particular that this also covers the cases where $m_C(q) < m_i^*(T_S)$, but this cannot occur because in that case G will not execute a coup at T_S , and if the solution to (7) is that small, R 's optimal choice is to optimize the “no-coup” scenario.

We conclude that the optimal payoff, $U_C(m_C(q))$, is strictly decreasing in q (it is clearly continuous).

Finally, we show that at $q = 1$, the ruler prefers to play the riskless strategy:

$$U_C(1) = \frac{\tilde{m}}{(1 + \theta\tilde{m})(\tilde{m} + T_S)} - \tilde{m} < \frac{\theta m_i^*(T_S)}{\theta m_i^*(T_S) + T_S} - m_i^*(T_S) = U_N(1),$$

where the inequality follows from

$$\frac{\tilde{m}}{(1 + \theta\tilde{m})(\tilde{m} + T_S)} - \tilde{m} < \frac{\tilde{m}}{\tilde{m} + T_S} - \tilde{m} < \frac{\theta\tilde{m}}{\theta\tilde{m} + T_S} - \tilde{m} \leq \frac{\theta m_i^*(T_S)}{\theta m_i^*(T_S) + T_S} - m_i^*(T_S),$$

where the last inequality follows from $m_i^*(T_S)$ being the optimizer under complete information.

We have now established that $U_C(0) > U_N(0)$, $U_C(1) < U_N(1)$, that U_N is strictly increasing while U_C is strictly decreasing. Since both functions are continuous, it follows that there exists precisely one intersection, at some $q^* \in (0, 1)$, such that R strictly prefers the risky strategy for all $q < q^*$, and strictly prefers the riskless one for all $q > q^*$.

Since $\theta > 1$ makes hiring a general strictly preferable to not hiring one as long as the probability of a coup is zero, it follows that with $\theta > 1$ R will always hire a general (if R prefers the risky strategy to the one that ensures that no coup takes place, then she must prefer it to not hiring G as well). Conversely, $\theta < 1$ ensures that R does not hire anyone.

The final claims of the lemma follow immediately: if $m_C(q) < m_i^*(T_L)$ when the risky strategy is chosen, the allocation obviously falls short of the optimum to deal with the large threat.¹⁰ Since $m_i^*(T_S) < m_i^*(T_L)$, the same is certainly true under the safe strategy. ■

Proof of Lemma 7. We establish this result by showing that both U_N and U_C are increasing in θ regardless of the value of q .

Consider U_N first and start with θ sufficiently small so that $\sqrt{\theta T_S} - T_S \leq S_i^*(T_S)$; that is, any θ that makes the complete-information constraint not binding against T_S

¹⁰For example, this happens when $b_i = 0.2$, $c = 0.3$, $\theta = 20$, $T_S = 1$, and $T_L = 7$. In this case $q^* \approx 0.055$, while $m_C(q) < m_i^*(T_L)$ for all $q > 0.005$.

so that $m_i^*(T_S) = \max(0, \sqrt{T_S/\theta} - T_S/\theta)$. If $\theta \leq T_S$, then $m_i^*(T_S) = 0$, and $U_N = 0$ for any such θ . (This means that R will rather go it alone than a hire a general even when doing so means no coup will occur.) If, on the other hand, $\theta > T_S$, then $\theta m_i^*(T_S) = \sqrt{\theta T_S} - T_S > 0$, so we can write

$$U_N = q \left(1 - \sqrt{\frac{T_S}{\theta}} \right) + (1 - q) \left(\frac{\sqrt{\theta T_S} - T_S}{\sqrt{\theta T_S} - T_S + T_L} \right) - \sqrt{\frac{T_S}{\theta}} + \frac{T_S}{\theta}.$$

Taking the derivative with respect to θ and setting it greater than zero yields, after some algebra,

$$q + (1 - q) \left[\frac{\theta T_L}{(\sqrt{\theta T_S} - T_S + T_L)^2} \right] + 1 > 2\sqrt{\frac{T_S}{\theta}}.$$

Since $\theta > T_S \Rightarrow \sqrt{T_S/\theta} < 1$, this inequality will hold whenever

$$q + (1 - q) \left[\frac{\theta T_L}{(\sqrt{\theta T_S} - T_S + T_L)^2} \right] > \sqrt{\frac{T_S}{\theta}}$$

obtains. But since the left-hand side is a linear combination of 1 and the bracketed term, the fact that $\sqrt{T_S/\theta} < 1$ further tells us that this inequality will hold whenever

$$\frac{\theta T_L}{(\sqrt{\theta T_S} - T_S + T_L)^2} > \sqrt{\frac{T_S}{\theta}},$$

obtains, which we can establish as follows. Taking the derivative of the left-hand side with respect to T_L yields

$$\frac{\theta (\sqrt{\theta T_S} - T_S + T_L)}{(\sqrt{\theta T_S} - T_S + T_L)^3} > 0,$$

and since this means that it is strictly increasing, it is sufficient to establish the inequality for the smallest value T_L can hold; that is, it is sufficient to establish the inequality for $T_L = T_S$. But in this case, the left-hand side reduces to 1, and we already know that $1 > \sqrt{T_S/\theta}$. Thus, we conclude that U_N is strictly increasing in θ whenever the optimal complete-information allocation is unconstrained and positive.

Consider now θ high enough so that $\sqrt{\theta T_S} - T_S > S_i^*(T_S)$; that is, any θ that makes the complete-information constraint binding against T_S so that $m_i^*(T_S) = S_i^*(T_S)/\theta$. Since $S_i^*(T_S)$ is constant in θ , the inequality will be preserved for any larger θ as well. But now we obtain $\theta m_i^*(T_S) = S_i^*(T_S)$, so we can write

$$U_N = \frac{q S_i^*(T_S)}{S_i^*(T_S) + T_S} + \frac{(1 - q) S_i^*(T_S)}{S_i^*(T_S) + T_L} - \frac{S_i^*(T_S)}{\theta}, \quad (9)$$

which is clearly increasing in θ . Thus, once θ is high enough that the complete-information constraint binds, increasing it further will only increase U_N as well (since the constraint will continue to bind).

Let us now establish the equivalent claim for U_C . We have two cases to consider.

Case 1: $m_C = m_i^*(T_L)$, which we recall from the proof of Lemma 6 further means that $m_C = S_i^*(T_L)/\theta$. Substituting this into (6) yields

$$U_C = \left(\frac{q}{1 + S_i^*(T_L)} \right) \left(\frac{S_i^*(T_L)}{S_i^*(T_L) + \theta T_S} \right) + \frac{(1 - q)S_i^*(T_L)}{S_i^*(T_L) + T_L} - \frac{S_i^*(T_L)}{\theta}, \quad (10)$$

from which we obtain

$$\frac{dU_C}{d\theta} = \frac{S_i^*(T_L)}{\theta^2} - \frac{qT_S S_i^*(T_L)}{(1 + S_i^*(T_L))(S_i^*(T_L) + \theta T_S)^2} > 0,$$

where the inequality can be established with simple algebra. Thus, U_C is strictly increasing in θ whenever m_C is the constrained solution.

Case 2: m_C is the unconstrained optimizer so the FOC is satisfied: $\frac{\partial U_C}{\partial m} = 0$ at the optimum. We can simply apply the envelope theorem to obtain

$$\frac{dU_C}{d\theta} = \frac{\partial U_C}{\partial m} \frac{dm}{d\theta} + \frac{\partial U_C}{\partial \theta} = \frac{\partial U_C}{\partial \theta} = m_C \left[\frac{(1 - q)T_L}{(\theta m_C + T_L)^2} - \frac{q m_C}{(m_C + T_S)(1 + \theta m_C)^2} \right],$$

which tells us that U_C must be increasing in θ if

$$\frac{(1 - q)T_L}{(\theta m_C + T_L)^2} > \frac{q m_C}{(m_C + T_S)(1 + \theta m_C)^2}. \quad (11)$$

Since (7) is satisfied, we know that

$$\frac{(1 - q)T_L}{(\theta m_C + T_L)^2} = \left(\frac{1}{\theta} \right) \left[1 - \frac{q(T_S - \theta m_C^2)}{(m_C + T_S)^2(1 + \theta m_C)^2} \right].$$

We substitute this into (11) and after some algebra reduce that inequality to

$$(m_C + T_S)^2(1 + \theta m_C) > qT_S,$$

which clearly holds: $(m_C + T_S)^2(1 + \theta m_C) > (m_C + T_S)^2 > m_C + T_S > T_S > qT_S$. Thus, if m_C is the unconstrained optimizer, U_C is strictly increasing in θ . ■

Proof of Lemma 8. We shall establish this result by showing that both U_N and U_C are either constant in b_i or strictly increasing.

We begin with U_N . If $m_i^*(T_S)$ is the unconstrained complete-information optimum, then it is independent of b_i , and so U_N itself is constant in b_i . If $m_i^*(T_S) =$

$S_i^*(T_S)/\bar{\theta}$, on the other hand, then the allocation is strictly increasing in b_i because $S_i^*(T_S)$ does. The payoff in this case is given by (9). Since

$$\frac{dU_N}{db_i} = \frac{\partial U_N}{\partial S_i^*} \frac{dS_i^*}{db_i} + \frac{\partial U_N}{\partial b_i}$$

but $\frac{\partial U_N}{\partial b_i} = 0$ and $\frac{dS_i^*}{db_i} > 0$, it follows that

$$\text{sgn} \left(\frac{dU_N}{db_i} \right) = \text{sgn} \left(\frac{\partial U_N}{\partial S_i^*} \right).$$

Thus, we need to show that

$$\frac{\partial U_N}{\partial S_i^*} = \frac{qT_S}{(S_i^*(T_S) + T_S)^2} + \frac{(1-q)T_L}{(S_i^*(T_S) + T_L)^2} - \frac{1}{\bar{\theta}} > 0.$$

We are going to split the proof in two cases. First, suppose that $S_i^*(T_S) < \sqrt{T_S T_L}$, which implies that $T_S (S_i^*(T_S) + T_L)^2 > T_L (S_i^*(T_S) + T_S)^2$. We can rewrite the condition on the derivative as

$$\bar{\theta} > \frac{(S_i^*(T_S) + T_S)^2 (S_i^*(T_S) + T_L)^2}{qT_S (S_i^*(T_S) + T_L)^2 + (1-q)T_L (S_i^*(T_S) + T_S)^2} \equiv \underline{\theta}.$$

By Assumption 3, $\bar{\theta} > (\sqrt{T_S} + \sqrt{T_L})^2$, so it suffices to show that $(\sqrt{T_S} + \sqrt{T_L})^2 > \underline{\theta}$. But since $S_i^*(T_S) < \sqrt{T_S T_L}$, it follows that

$$\underline{\theta} < \frac{(S_i^*(T_S) + T_S)^2 (S_i^*(T_S) + T_L)^2}{T_L (S_i^*(T_S) + T_S)^2} = \frac{(S_i^*(T_S) + T_L)^2}{T_L},$$

so we only need to show that

$$\left(\sqrt{T_S} + \sqrt{T_L} \right)^2 > \frac{(S_i^*(T_S) + T_L)^2}{T_L} \Leftrightarrow S_i^*(T_S) < \sqrt{T_S T_L}.$$

Since the last inequality is true by supposition, the claim holds.

Turning now to the other possibility, suppose that $S_i^*(T_S) > \sqrt{T_S T_L}$, which implies that $T_S (S_i^*(T_S) + T_L)^2 < T_L (S_i^*(T_S) + T_S)^2$. Recall that $m_i^*(T_S)$ is the binding allocation, which means that $S(T_S) > S_i^*(T_S)$, which implies that

$$\bar{\theta} > \frac{(S_i^*(T_S) + T_S)^2}{T_S}.$$

But this now means that

$$\frac{\partial U_N}{\partial S_i^*} > \frac{(1-q)T_L}{(S_i^*(T_S) + T_L)^2} - \frac{(1-q)T_S}{(S_i^*(T_S) + T_S)^2},$$

so it suffices to show that

$$\frac{T_L}{(S_i^*(T_S) + T_L)^2} > \frac{T_S}{(S_i^*(T_S) + T_S)^2} \Leftrightarrow S_i^*(T_S) > \sqrt{T_S T_L}.$$

Since the last inequality is true by supposition, the claim holds. Thus, U_N is non-decreasing in b_i .

Consider now U_C . If m_C is the unconstrained optimizer, then $\left. \frac{\partial U_C}{\partial m} \right|_{m_C} = 0$. The envelope theorem then tells us that

$$\frac{d U_C}{d b_i} = \left. \frac{\partial U_C}{\partial m} \frac{d m}{d b_i} \right|_{m_C} + \frac{\partial U_C}{\partial b_i} = \frac{\partial U_C}{\partial b_i} = 0,$$

which means that U_C is independent of b_i in this case.

If, on the other hand, m_C is the constrained optimizer, then $\left. \frac{\partial U_C}{\partial m} \right|_{m_C} > 0$ and $m_C = m_i^*(T_L) = S_i^*(T_L)/\theta$. Since $\left. \frac{\partial U_C}{\partial b_i} \right|_{m_C} = 0$, we obtain

$$\frac{d U_C}{d b_i} = \left. \frac{\partial U_C}{\partial m} \frac{d m}{d b_i} \right|_{m_C} > 0,$$

where the inequality follows from $\left. \frac{d m}{d b_i} \right|_{m_C} = \left(\frac{d S_i^*}{d b_i} \right) / \theta > 0$ and $\left. \frac{\partial U_C}{\partial m} \right|_{m_C} > 0$. Thus, U_C is non-decreasing in b_i as well.

We conclude that the payoffs are strictly increasing whenever $m_i^*(T_S) = S_i^*(T_S)/\bar{\theta}$ (in the riskless subgame) or $m_i^*(T_L) = S_i^*(T_L)/\bar{\theta}$ (in the risky subgame) are the constrained optima under complete information.

Recall that S_i^* itself is increasing in b_i , that $S(T)$ is independent of b_i , that $S_i^*(T_S) < S_i^*(T_L)$, and that $S(T_S) < S(T_L)$ under Assumption 3. Consider now very low values of b_i (and possibly c) such that the loyalty constraint binds in both cases: $S_i^*(T) < S(T)$ for $T \in \{T_S, T_L\}$. In other words, consider $b_i < b_1$.¹¹ The results above indicate that R 's payoff from both U_N or U_C is strictly increasing in b_i , so she must pick the highest such b_i that still ensures that the constraints obtain. If $\bar{b} \leq b_1$, then R must select from the most privileged group regardless of q .

If $\bar{b} \in (b_1, b_2)$, then at least one of the constraints will cease to be binding. The corresponding payoff will now be constant in b_i whereas the other one will continue to increase. If $b_1 = b^*(T_S)$, then the constraint that affects U_N will no longer bind. R is now indifferent among any $b_i \in [b_1, \bar{b}]$ when the equilibrium outcome is riskless, which we know to be the case for any $q > q^*$. On the other hand, since $b_1 = b^*(T_S)$ implies that $b_2 = b^*(T_L)$, it follows that $\bar{b} < b^*(T_L)$, so

¹¹ $b^*(T)$ can be concave or strictly decreasing in T , depending on the values of c , which is why we cannot say which constraint will be relaxed first in general.

the constraint is still binding for the risky continuation game. Since U_C is strictly increasing in b_i , R must strictly prefer to pick \bar{b} for any $q \leq q^*$. The situation where $b_1 = b^*(T_L)$ is analogous, *mutatis mutandis*.

Finally, if $\bar{b} \geq b_2$, then the constraints are not binding in either continuation game, so R must be indifferent among any $b_i \in [b_2, \bar{b}]$ regardless of q . ■