Chapter 8

Resolving the Monitoring Problem: The Hidden Value of "Worthless" Weapons Inspections

On December 16, 2003, ABC News' Diane Sawyer sat down for an interview with President George W. Bush. Nine months had passed since the start of the Iraq War, and the U.S. had failed discover any evidence of Iraqi weapons of mass destruction. Pressed on this point, Bush emphasized that he acted on sound intelligence reports indicating Saddam Hussein's desire to obtain such weapons. Sawyer noted the distinction between intentions and realities, Bush rhetorically asked "So what's the difference?"[1]

Bush's emphasis on intentions over realities highlights an important problem for the nonproliferation regime. Throughout this book, we have assumed that the declining state has perfect knowledge of the rising state's arming decision. As a result, the declining state could leverage the threat of preventive war to force the rising state not to proliferate without having to offer inducements.

However, nuclear programs are often clandestine. The corresponding uncertainty gives pause to a declining state for two reasons. First, preventive war threats alone can no longer convince a rising state not to build. After all, if the rising state is dissatisfied with the status quo, it could proliferate

[1] Quoted in Isikoff and Corn 2006, 342.

in secret and reap the benefits later on. Second, butter-for-bombs inducements lack visual confirmation; if a declining state tries buying off its rival, it must have faith that the rival will not take the concessions and proliferate anyway. Declining states may therefore be pessimistic that butter-for-bombs agreements can work at all in such an environment.

To understand how hidden actions impact the bargaining environment, this chapter extends Chapter 3's model to allow for imperfect information. Whereas the declining state could perfectly observe the rising state's investment decision previously, it now remains the dark. Consequently, logical inference–not visual confirmation–must guide the declining state's decision whether to sit and wait or launch preventive war.

Unable to monitor the rising state's investment decision, the declining state can take one of two strategic approaches. First, it could outright ignore its informational issues and naively continue bargaining as it has previously. Second, it could resize its offers to proactively combat the problem. Unsurprisingly, the declining state fares better with the latter approach.

Specifically, with imperfect information, the interaction leads to four types of behavior. When the extent of the power shift is small, the declining state infers that rejecting reasonable offers strictly dominates investing in weapons. In turn, the declining state outright ignores the rising state's ability to build and offers the rising state its reservation value for war. Without a better alternative, the rising state accepts.

When the extent of the power shift is still small but worth the cost of proliferation, the states engage in the familiar butter-for-bombs deal; the declining state offers immediate concessions, while the rising state accepts the concessions and does not build. Perhaps surprisingly, imperfect information adds no further complication to the interaction. Indeed, the optimal size of butter-for-bombs offers manipulates the rising state's opportunity cost for building in such a way that the rising state simply finds building unprofitable. As such, the declining state rests assured that the rising state will not build despite no direct oversight.

The declining state only faces a dilemma when the extent of the power shift is great. With perfect information, the declining state could leverage the threat of preventive war to force the rising state not to build. But this is no longer an option with imperfect information. Now the declining state must choose whether to lowball the rising state or to continue to provide butter-for-bombs deals.

Lowballing puts the states in an awkward situation. If the declining state does not prevent, the rising state prefers building; if the rising state builds, the declining state prefers preventing; if the declining state prevents, the rising state prefers not building so as not to waste the investment cost; and if the rising state does not build, the declining state prefers not preventing. This reaction cycle causes the parties to randomize. However, the rising state's optimal mixed strategy holds the declining state to its pre-shift war payoff. Thus, as long as the power shift is not too extreme, the declining state opts to offer a butter-for-bombs deal to steal some of the surplus. And once again, the structure of butter-for-bombs agreements ensures that the rising state has no incentive to build, thus alleviating the declining state's proliferation concerns.

On the other hand, in the case of an extreme power shift, the declining state cannot improve over its war payoff. As such, in the best case scenario, this triggers the mixing behavior described above. Proliferation occurs with positive probability here, as does preventive war. Unfortunately, this implies that sometimes the declining state initiates a preventive war even though the rising state did not invest in weapons.

Given how disastrous the final outcome is, we may wonder whether states could find an institutional way out. Similar to Chapter 5's exploration of endogenous investment costs, we show that using institutions to increase the cost of nuclear weapons leads to a Pareto-improving equilibrium outcome. Here, increasing the rising state's burden convinces the declining state to offer butter-for-bombs deals. In turn, the rising state can capture more of the surplus. Surprisingly, *potential proliferators benefit from facing greater barriers to proliferation.*

We may then wonder what sort of institutions build such obstacles. While Chapter 5 focused normative methods, this chapter's focus on hidden actions pushes the discussion toward weapons inspections. Interestingly, weapons inspectors can be extremely valuable while appearing completely worthless. By investigating likely weapons sites, inspectors cut off rising states' cheapest avenues to proliferation. This has the aforementioned positive effect of making nonproliferation deals credible. Yet this means that states planning to violate agreements never admit inspectors, while states that admit inspectors have no violations to be found. Consequently, inspectors increase welfare while simultaneously appearing to serve no purpose.

This chapter proceeds as follows. The next section reviews a related model from Debs and Monteiro (2013), highlighting key differences in our assumptions. The section after contains a modified version of Chapter 3's model to incorporate imperfect information; negotiating over the weapons leads to efficient outcomes in most cases. The fourth section shows that institutions can improve both sides' welfare in the rare case where negotiations fail; surprisingly, the rising state benefits by having a more difficult time constructing weapons. We then use that institutional logic to argue that imperfect information did *not* cause the Iraq War. A conclusion follows.

8.1 Known Unknowns, Weapons Inspections, and the Cost of Inconvenience

In a recent article, Debs and Monteiro (2013) construct a related model of proliferation and preventive war deterrence. While their model deemphasizes negotiating over the rising state's weapons program, they show that preventive war occurs with positive probability if preventive war is cheap and the effect of a power shift is great. The result is novel and intriguing but logically demanding. As such, we begin by untangling their model's complex strategic behavior.

To begin, consider both the rising state's decision to build and the declining state's decision to prevent. In a world with imperfect information, these decisions are effectively simultaneous. Figure 8.1 illustrates the potential outcomes of such a game. If the declining state prevents, both sides receive war payoffs, with the rising state also losing its investment cost if it built in the meantime. Inaction from both parties results in the status quo, while weapons construction without preventive war leads to a successful power shift.

Suppose the result of a successful power shift is worth the investment and the cost of preventive war is small. What is the result of such a simultaneous move interaction? Debs and Monteiro show that both parties mix. The declining state sometimes prevents and sometimes does not. The rising state sometimes builds and sometimes does not. All outcomes occur with positive probability.

To see why, note that none of the outcomes can occur with certainty. If the rising state is building, the declining state would respond by preventing. But if the declining state is preventing, then the rising state would not build so as to avoid the wasted investment cost. Yet if the rising state is not building, the declining state would maintain the status quo by not preventing. And

	Build	Pass
Prevent	**Preventive War, Wasted Costs**	**Preventive War**
Advance	**Successful Power Shift**	**Status Quo**

Figure 8.1: A simultaneous move game between a declining state (rows) and rising state (columns). If preventive war and the cost of proliferation are cheap, the actors must both mix in equilibrium.

if the declining state is not preventing, the rising state ought to build. The cycle begins anew. Therefore, the states must optimally randomize.

Debs and Monteiro show that this logic holds even if the declining state receives a noisy about the rising state's decision, provided that the signal is sufficiently weak. They then use this logic to explain the outbreak of the Iraq War. During the lead up, the Bush administration emphasized that the United States could not effectively monitor the situation on the ground, as weapons inspections were an imperfect source of information revelation. As such, preventive war was a rational decision–even if the U.S. discovered *ex post* that Iraq lacked a serious weapons program.

The model and subsequent analysis in this chapter break from Debs and Monteiro in three key respects. First, consistent with the theme of this book, we focus on the declining state's decision to offer concessions to the rising state upfront so as to reduce the need to invest in weapons. Differing from previous chapters, though, the declining state is unable to directly observe whether its bribe worked at the time it chooses whether to prevent.

Second, we consider how endogenous intelligence gathering and inspection institutions alter the bargaining landscape in the shadow of preventive action. Despite asymmetric information's lauded status as a rationalist explanation for war, the theoretical literature on how states might resolve information problems short of fighting remains surprisingly barren. The lone exception–Arena and Wolford (2012)–focuses on intelligence as an information gathering device. They allow a state to endogenously choose a level of military capability and intelligence subject to the same budget constraint. The intelligence investment then yields a noisy signal about a second state's type, which in turn allows the first state to better calibrate its offer. Thus, their focus is on resolving *incomplete* information as a cause of war.[2]

However, our focus here is on *imperfect* information. Research on United Nations peacekeeping has identified that a primary function of U.N. missions is to reveal hidden actions–whether in military buildups (Fortna 2004, 485-489) or electoral violations (Weiss, Forsythe, and Coate 1997, 53-54). In a similar vein, weapons inspectors in the Debs and Monteiro model provide intelligence through the exogenously given strength of the noisy signal. And, indeed, the search for information guided UNMOVIC. According to UNSC Resolution 1284, which established UNMOVIC, the purpose of the organization was to

> establish and operate...a reinforced system of ongoing monitoring and verification, which will implement the plan approved by the Council in resolution 715 (1991) and address unresolved disarmament issues, and that UNMOVIC will identify, as necessary in accordance with its mandate, additional sites in Iraq to be covered by the reinforced system of ongoing monitoring and verification.

In contrast, we stress that weapons inspections alter the cost of proliferation–such inspections shut down the most effective avenues to proliferate so long as the declining state would intervene if inspectors discovered active weapons programs. As the next two sections show, this leads to Pareto improving agreements.

Third, we differentiate between uncertainty from imperfect information and endogenous uncertainty. Debs and Monteiro (19) claim that "the problem was that in order to change U.S. policy, intelligence reports would have to prove that Iraq *did not have and would not develop WMD*." The United States being unable to directly observe Iraq's previous actions equates to standard imperfect information, which is common in simultaneous move games such as the one in Figure 8.2.

However, rational choice theory maintains that actors need not directly see a competitor's action to make a well-informed response. Note that up strictly dominates down for player 1 in Figure 8.2. In turn, player 2 can infer that player 1 will not play down. This inference does not come from direct observation. Indeed, player 2 will not know what player 1 has chosen

[2] Similarly, work on international mediation shows that a mediator can resolve these problems by revealing private information to the potential warring parties (Kydd 2003; Rauchhaus 2006).

	Left	Right
Up	1, 1	0, -1
Down	-1, -10	-3, 0

Figure 8.2: A simple simultaneous move game. Down is not a rationalizable strategy. Thus, player 2 anticipates player 1's move up and responds by choosing left.

until she makes a move for herself. Yet player 2 can anticipate player 1's up strategy simply because down is not rationalizable. Thus, while the game has uncertainty from imperfect information, the game has no endogenous uncertainty–the players know exactly which move the other side will pick. As a result, player 2 can safely pick left even though her outcome is disastrous if she infers incorrectly.

We will see similar logic throughout the modeling section. When the declining state properly incentivizes the rising state, proliferation becomes a dominated strategy. Thus, despite the uncertainty from imperfect information, there is no endogenous uncertainty. The declining state knows that the rising state will not randomly invest in weapons and properly responds by keeping the peace. Thus, with this incentive-compatible induced trust, commitment problems do not occur.

8.2 Bargaining with Imperfect Monitoring

Per usual, we extend Chapter 3's model, this time by integrating imperfect information. As with last chapter, since we have already established that butter-for-bombs agreements are sustainable in the long-term, we focus on a two-period interaction. In the pre-shift state of the world, D offers $x \in [0, 1]$. R rejects, accepts, or builds. Rejecting ends the interaction and gives each side its pre-shift war payoff. Accepting means that R forgoes having nuclear weapons if the interaction reaches the post-shift period, whereas building incurs the cost $k > 0$ but ensures that R will increase its power in the post-shift period.

Unlike previous versions, however, D cannot differentiate between R accepting and R building.[3] At this point, D must prevent or pass. Preventing ends the game in pre-shift war. If R accepted, passing ends the interaction and locks in the offered distribution for the rest of time. If R built, then the game advances to the post-shift phase. D now sees that R built and offers $y \in [0,1]$.[4] R accepts or rejects. Accepting ends the game with that distribution for the rest of time, while rejecting gives each side its post-shift war payoff.

Throughout our exploration of the monitoring problem, we assume that $k > \delta(p'_R - c_R)$ for two reasons. First, our goal is to understand the origins of proliferation. But when k falls below that critical value, Chapter 4 already showed that bargaining will fail to yield an agreement. Second, Chapter 5 argued that Pareto-improving institutions ought to endogenously inflate k to be that large. Therefore, if a new mechanism causes proliferation to occur here, it must be that $k > \delta(p'_R - c_R)$.

The following subsections give the results of the interaction. To preview, Figure 8.3 plots the outcomes as a function of the cost of proliferation and extent of the power shift. It has many similarities to Figure 3.5. If the cost of weapons is too great relative to the power shift, no proliferation occurs; note that this range now extends to regions where the threat of preventive war had previously deterred proliferation. In the middle range, butter-for-bombs agreements work just as they had before. In the range extending up and to the right, butter-for-bombs agreements remain effective, though the D must offer additional concessions to R. Preventive war and proliferation occur only in the top left region.

8.2.1 When Monitoring Problems Are Irrelevant

Since this is a sequential game of complete but imperfect information, we seek its subgame perfect equilibria. This may seem counterintuitive first, since

[3]Thus, unlike Debs and Monteiro's model, this interaction has no noisy signal. This is a "worst-case scenario" analysis–we show that even with no information, institutions can lead to Pareto improvement. Reaching the Pareto frontier would only be easier with a noisy signal.

[4]Intuitively, this means D is unable to see technological progress until R can credibly demonstrate it. For example, upon proliferating, R might publicly test a nuclear weapon to demonstrate its strength to D, thereby alleviating the information problem. See Meirowitz and Sartori 2008 and Reed, Wolford, and Arena 2013 for models when this is not the case.

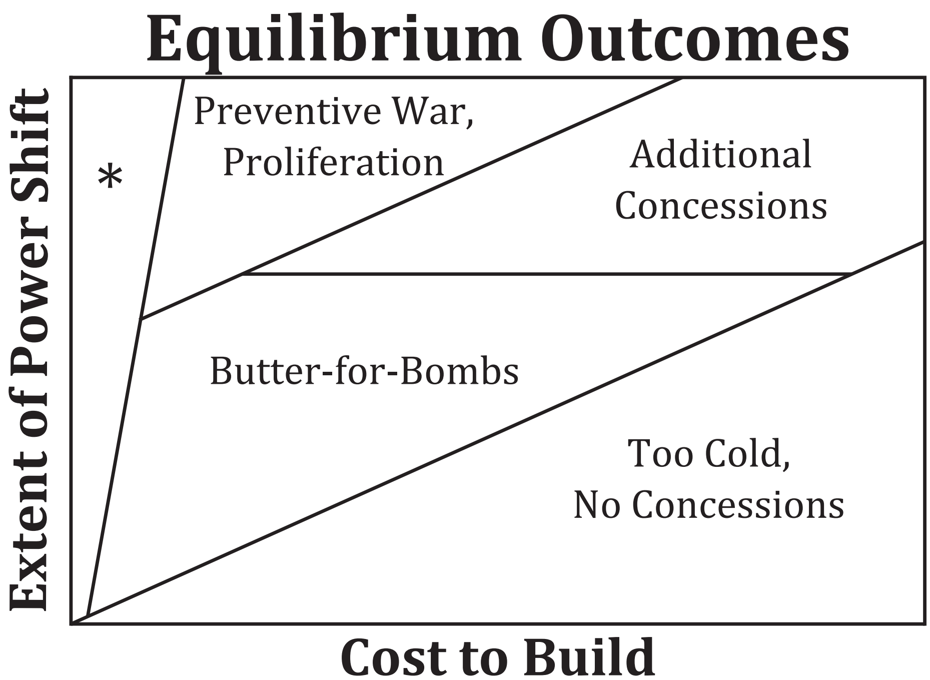

Figure 8.3: Equilibrium outcomes as a function of the cost to build and the rising state's level of future power. See Chapter 5 for a discussion of the region marked with the asterisk.

the monitoring problem appears to create incomplete information. However, the inability to monitor R simply means that D must make its decision to prevent simultaneous to R's decision to build. Thus, the game only has *imperfect* information. Since optimal strategies do not depend on prior beliefs, we therefore search for subgame perfect equilibria and not perfect Bayesian equilibria.

Because the post-shift state of the world is identical to the game with perfect monitoring, Lemma 3.1 still applies. In every post-shift period, D offers $y = p'_R - c_R$ and R accepts. D earns $1 - p'_R + c_R$ and R earns $p'_R - c_R$ for the rest of time. We pick up on the search for stationary SPE from here, starting with situations in which the investment fails to bring sufficient rewards:

Proposition 8.1. *If $p'_R - p_R < \frac{k(1-\delta)}{\delta}$, D offers $x = p_R - c_R$ and R accepts in every SPE.*[5]

Intuitively, Proposition 8.1 is analogous to Proposition 3.2.[6] If the power shift is too small relative the cost of investment, R prefers to reject if the concessions are too small (below $p_R - c_R$) and prefers to accept otherwise. Faced with this, D prefers offering R's reservation value for war and enjoying the entire surplus. Note that the monitoring problem plays no role here at all; D knows R will not build because the investment is unreasonably costly.

8.2.2 Why Butter-for-Bombs Agreements Still Work

Proposition 8.1 should not come as a surprise, since we have seen similar results in previous chapters. But what happens when proliferation is a credible threat? Initially, we might think that all butter-for-bombs agreements will immediately fail as a result of the monitoring problem–D cannot monitor R's compliance. Proposition 8.2 shows that this fear is overstated:

Proposition 8.2. *If $p'_R - p_R \in (\frac{k(1-\delta)}{\delta}, \frac{c_D+c_R}{\delta})$, D offers $x = p'_R - c_R - \frac{k(1-\delta)}{\delta}$ and R accepts in every SPE.*

[5] The equilibria in this chapter are not unique. Off the equilibrium path of play, the simultaneous move portion of the interaction has multiple equilibria if $x = p_R + c_D$. But since this off the path in all but the parameter range for Proposition 8.4, the equilibrium outcomes are unique. See Lemma 8.3 in this chapter's appendix for an explanation.

[6] Its proof is nearly identical as well and therefore omitted.

Note that Proposition 8.2 generates the same result as Proposition 3.3–that is, butter-for-bombs agreements that exist in a world with perfect monitoring also exist in a world with completely imperfect monitoring.

Why is D's lack of information irrelevant? Consider R's incentives. When it receives $x = p'_R - c_R - \frac{k(1-\delta)}{\delta}$, R does not choose to accept out of fear that D will prevent. Rather, R opts against building because D's offer is so attractive that additional weapons are no longer profitable; the size of the offer endogenously manipulates R's opportunity cost to ensure this. As a result, even if D is completely blind to R's actions, it can still rest assured that R will not break the butter-for-bombs agreement. The appendix contains proof to verify this intuition.

8.2.3 The Monitoring Problem Premium

Outside of Chapter 3's butter-for-bombs region, proliferation seems unavoidable. But this intuition is wrong once again. Proposition 8.3 shows that the monitoring problem causes D to expand the circumstances under which it is willing to make butter-for-bombs concessions. Again, the states avoid the proliferation outcome:

Proposition 8.3. *If $p'_R - p_R \in (\frac{c_D+c_R}{\delta}, c_D + c_R + \frac{k(1-\delta)}{\delta})$, D offers $x = p'_R - c_R - \frac{k(1-\delta)}{\delta}$ and R accepts in every SPE.*

Why does the monitoring problem cause D to become more amenable to making a deal? To understand the logic, we must consider the off-the-equilibrium path consequences of being less magnanimous with concessions. Recall that R requires at least $x = p'_R - c_R - \frac{k(1-\delta)}{\delta}$ to forgo investment. Figure 8.4 illustrates the simultaneous move subgame the states face. Suppose D offers some amount less than that. Since $p'_R - p_R > \frac{c_D+c_R}{\delta}$, D prefers to prevent if R builds but prefers passing if R does not build. Meanwhile, R has contrary incentives; it wants to build if D passes and accept (and avoid wasting the investment cost) if D prevents.

Note that these strictly opposing preferences imply that no pure strategy Nash equilibrium exists for the subgame. But since this is a finite game, Nash's (1951) theorem ensures that an equilibrium must exist, so we look in mixed strategies. If that equilibrium is in totally mixed strategies–and the proof in this chapter's appendix shows that it is–then the indifference conditions require D's expected utility for preventing to be equal to its expected utility for passing. But note that regardless of R's strategy, D earns

	Build	Pass
Prevent	**1-p_R-c_D, p_R-c_R-k(1-δ)**	**1-p_R-c_D, p_R-c_R**
Advance	**(1-x_t)(1-δ)+δ(1-p'$_R$+c_R), x_t(1-δ)+δ(p'$_R$-c_R)-k(1-δ)**	**1-x_t, x_t**

Figure 8.4: The simultaneous move decision. D's choices are the rows; R's choices are the columns.

$1 - p_R - c_D$ for preventing. Thus, in the mixed strategy Nash equilibrium, D must earn $1 - p_R - c_D$.

All told, mixing leads to a terrible outcome. Inefficiency abounds. The states fight war with positive probability, ensuring D receives its war payoff and none of the surplus. R sometimes pays the cost of investment. And worst of all, the outcome permits D to sometimes fight an accidental preventive war–that is, the states sometimes reach the outcome in which D prevents despite the fact that R did not build.

In contrast, D could offer the standard butter-for-bombs amount, $x = p'_R - c_R - \frac{k(1-\delta)}{\delta}$. As the intuition for Proposition 8.2 explained, this prompts R to accept regardless of the monitoring problem. In turn, D receives the remainder, or $1 - p'_R + c_R + \frac{k(1-\delta)}{\delta}$. Given that offering a smaller amount leads to the inefficient mixing, D prefers making the optimal butter-for-bombs offer to receiving its pre-shift war payoff if:

$$1 - p'_R + c_R + \frac{k(1-\delta)}{\delta} > 1 - p_R - c_D$$

$$p'_R - p_R < c_D + c_R + \frac{k(1-\delta)}{\delta}$$

This holds for the Proposition 8.3's parameters. Therefore, in the absence of effective monitoring, D prefers extending butter-for-bombs offers in these cases.

Why does trust still succeed, despite D's desire to prevent if R were to proliferate? While D has great reason to distrust R if its offer is small, accepting $x = p'_R - c_R - \frac{k(1-\delta)}{\delta}$ is in R's rational self-interest. Thus, even though a disaster seems likely, D can build a working trust with R, knowing that its large offer removes the impetus for preventive war.

Again, this chapter's appendix contains a complete proof.

8.2.4 The Extreme: Preventive War and Proliferation

We now consider the extreme case, in which D prefers fighting to giving the concessions necessary to induce R not to proliferate. As Proposition 8.4 details, D cannot improve over its war payoff here:

Proposition 8.4. *If $p'_R - p_R > c_D + c_R + \frac{k(1-\delta)}{\delta}$, D receives its war payoff in every SPE. Preventive war, accidental preventive war, and successful proliferation are supported in SPE.*

Why are there multiple subgame perfect equilibria? Note that if D offers more than $p'_R - c_R - \frac{k(1-\delta)}{\delta}$, it prefers fighting regardless of R's decision to build. But if it offers an amount less than or equal to $p'_R - c_R - \frac{k(1-\delta)}{\delta}$, R prefers building if D does not prevent. This causes the mixing behavior from before, which guarantees that D earns $1 - p_R - c_D$. Yet D earns that amount if it offers more than $p'_R - c_R - \frac{k(1-\delta)}{\delta}$. Thus, D is indifferent to *all* of its offer sizes. Since D's offer is the first move of the interaction, SPE exist for any given offer size. Regardless, the result is inefficient: a large offer guarantees war with probability 1, while a smaller offer sees positive probability of war and investment.

Once more, the appendix contains a complete proof.

8.3 The Value of Weapons Inspections

The previous section showed the resolution to the nuclear proliferation dilemma depends heavily on the extent of the power shift. Although the results are efficient for most of the parameter space, inefficiency dominates when the power shift is great.

However, as of yet, we do not have a full understanding of the distributional consequences. Thus, this section analyzes the states' welfare and shows that both benefit from raising the investment cost when in the the parameters of Proposition 8.4 hold. We then show that similar welfare improvement occurs in a couple alternative specifications of the model.

8.3.1 Welfare

We begin with Theorem 8.1, which claims that both sides benefit from greater proliferation costs under certain conditions:

Theorem 8.1. *There exist greater values of k for which the equilibrium payoffs of the game strictly Pareto dominate the payoffs of all equilibria for when $k < \frac{\delta(p'_R - p_R - c_D - c_R)}{1-\delta}$.*

Note that the cost of proliferation determines whether the declining state is willing to bribe the rising state. If the cost of proliferation is extremely large, the declining state only needs to offer the rising state a minimal bribe and would thus prefer taking the remainder to fighting a war. But if the cost of proliferation is extremely cheap, the declining state prefers its war payoff and is unwilling to bribe.

In between these extremes, a particular cost of proliferation makes the declining state indifferent between bribing the rising state and fighting. Thus, at that cost, the declining state is willing to take the efficient route. Here, the declining state still receives its war payoff. The rising state, meanwhile, receives the entire surplus, which is more than what it would receive if costs were lower. Thus, *the rising state benefits from greater proliferation costs*, though the rising state's payoff decreases past this critical investment cost.

Figure 8.5 illustrates both states' payoffs as k increases, holding the other parameters at $p_R = .2$, $p'_R = .7$, $c_D = .25$, $c_R = .1$, and $\delta = .9$. When costs are high, D treats the bargaining problem as though power were static and extracts the entire surplus. When costs fall in the middle range, D makes butter-for-bombs concessions, inducing R not to build despite the monitoring problem. When costs are low, multiple equilibria exist. The payoffs shown are for the equilibrium in which D offers $x > p_R + c_D$. Like all other equilibria for the parameter range, this equilibrium is inefficient. The deadweight loss ensures that moving to the middle range will increase both sides' payoffs.[7]

8.3.2 Alternative Model Specifications

Before moving on, it is worth checking the robustness of the welfare results. Like all other models, this model is a tradeoff between simplicity and empirical validity. As such, we briefly consider three possible alternative specifications and explain the conditions under which the same Pareto improvement would exist.

[7] If we endogenized the cost of proliferation as in Chapter 5, R sets $k = \frac{\delta(p'_R - p_R - c_D - c_R)}{1-\delta}$ in the unique equilibrium. This value maximizes R's payoff in Figure 8.5 and means that R would intentionally inflate its investment cost even if it had unilateral control of the decision.

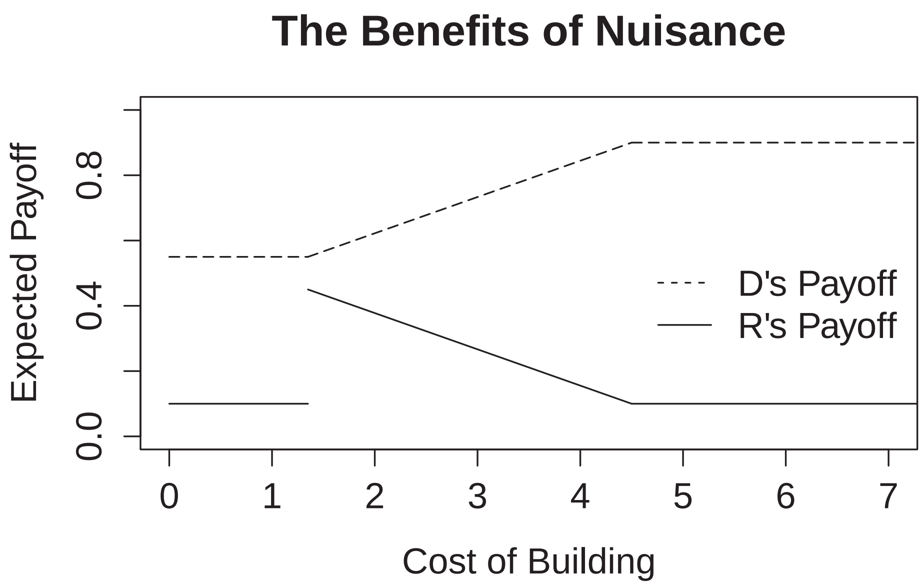

Figure 8.5: Equilibrium payoffs for both actors as a function of k. The outcomes associated with middling values of k Pareto dominate the outcomes associated with small values.

Infinite Horizon. For simplicity, this chapter investigated a two-period interaction. However, it is easy to verify that the same welfare improvement mechanism works if the states bargained repeatedly in the imperfect information equivalent of Chapter 3's model.

To see why, note that the strategies from Propositions 8.1, 8.2, and 8.3 remain an equilibrium in the infinite setup. In each of these cases, R does not build because it is not profitable to do so, and D knows not to prevent through iterated elimination of dominated strategies. Most importantly, this implies that R receives $p'_R - c_R - \frac{k(1-\delta)}{\delta}$ and D receives $1 - p'_R + c_R + \frac{k(1-\delta)}{\delta}$ if the parameters cross the $p'_R - p_R < c_D + c_R + \frac{k(1-\delta)}{\delta}$ threshold. Note that these payoffs are on the Pareto frontier.

Now consider the range in which $p'_R - p_R > c_D + c_R + \frac{k(1-\delta)}{\delta}$.For R and D to avoid proliferation or preventive war, R's continuation value must be at least $p'_R - c_R - \frac{k(1-\delta)}{\delta}$ and D's continuation value must be at least $p_R - c_D$; otherwise, R would prefer building or D would prefer preventing. However, $p'_R - c_R - \frac{k(1-\delta)}{\delta} + p_R - c_D > 1$ in this parameter space. The sum of the necessary continuation values is therefore more than the total pie the parties can divide. Thus, the actors cannot maintain the status quo in equilibrium.

In turn, D must prevent with positive probability or R must build with positive probability. Consider the first case. Due to the same indifference conditions as in the proof for Proposition 8.4 and Theorem 8.1, D must receive its war payoff in expectation, while R must receive strictly less than $p_R + c_D$ due to war's inefficiency. And also for the same reasons as Theorem 8.1, the parties would see Pareto improvement by inflating k into the range in which $p'_R - p_R < c_D + c_R + \frac{k(1-\delta)}{\delta}$.

Meanwhile, the second case cannot hold while the first case does not. If R builds with certainty, the proof for Proposition 8.4 shows that D prevents. We covered this case in the previous paragraph. Thus, the only remaining case is if R mixes. To be indifferent, D must offer an amount greater than or equal to $p'_R - c_R - \frac{k(1-\delta)}{\delta}$. But then D prefers preventing again, thus putting the actors back into the previous paragraph's problem. So D must prevent with positive probability, in turn assuring that inflating k yields Pareto improvement.

Costly Weapons Inspections. Weapons inspectors are not free–someone must pay for their employment, administrative support, and logistics. As such, inspections can only lead to welfare improvement if their cost is less than the inefficiency from warfare and weapons construction.[8]

Unfortunately, the question of whether inspections are cheaper than discord is a substantive one that we should evaluate on a case-by-case basis. Certainly, however, given the costliness of some wars and weapons programs, many situations easily allow for Pareto improving institutions. For example, UNMOVIC consisted of only 45 core members; both sides' militaries numbered in the hundreds of thousands in contrast.[9] Even in marginal cases, third party states have incentive to contribute to avoid war's negative externalities. And, indeed, Article XII of the International Atomic Energy Agency's charter and Article III of the NPT specifically build a weapons inspection regime.

Inspections as a Power Shift. Some rising states might be reluctant to let weapons inspectors into their countries, as information leaks could give declining states tactical advantages.[10] This is, in part, why declining states delegate weapons inspections to international organization. But international organizations are not perfect. To wit, the United Nations Special Commission on Iraq, the regime tasked with monitoring Iraqi compliance following the Persian Gulf War, ceased operations in 1999 amid allegations that Western intelligence agencies had infiltrated it (Blix 2004, 36-37; ElBaradei 2011, 32-33).

Fortunately, espionage does not doom inspection regimes. Like the case with costly weapons inspections, raising costs can still deliver Pareto improvement as long as the power shift is sufficiently small. Moreover, recall from Figure 8.5 that the rising state's payoff has a sizeable discontinuous increase when k crosses over the critical threshold. As such, rising states would accept a power shift to obtain the benefit from the discontinuous gain. Declining states, of course, would be happy to shift power in their favor.

[8] This is essentially a costly peace argument (Powell 2006; Coe 2012).

[9] On employment figures, see http://www.un.org/Depts/unmovic/documents/515en.PDF, page 2. More inspectors would have been necessary had negotiations not broken down in March 2003 (Blix 2003), but the total number would have still been exponentially smaller than the troops committed to combat. For further perspective, the yearly budget for IAEA inspections is about $120 million (ElBaradei 2011, 80).

[10] Indeed, for this precise reason, the credibility of the preventive war threat is necessary for the rising state to accept inspections. North Korea, for example, banished weapons inspections when it revamped its nuclear program. The United States failed to intervene. Had the United States been in better position to attack North Korea, it is unlikely the weapons inspectors would have left.

8.4 The Hidden Purpose (and Apparent Futility) of Weapons Inspections

How do weapons inspections alter the bargaining environment? The straightforward interpretation is that they provide information to the declining state. Suppose rising states cannot effectively hide their programs from the prying eyes of inspectors. If the declining state would launch preventive war short of a signal not to, the rising state might wish to invite weapons inspectors into its country and allay the declining state's fears. Since refusal to admit inspectors signals violations, the declining state can efficiently sort out the proliferators from the non-proliferators.

However, weapons inspectors are an imperfect solution. Proliferators have a home field advantage–violations could be anywhere in the country, leaving weapons inspectors with a lot of ground to cover without sufficient manpower. Strong intelligence alleviates some of the problem by giving inspectors likely locations of infractions, but that can be lacking as well since the absence of a smoking gun is not direct proof of compliance. In the extreme, cat-and-mouse games with inspectors might prevent inspections from providing any relevant informational content whatsoever.

Fortunately, even such ineffective weapons inspections have a hidden secondary effect that solves the proliferation problem. Although intelligence can be imperfect, the most efficient locations to construct weapons are evident. Weapons inspections effectively shut down these avenues to proliferation–inspectors can investigate such known sites and report violations. In turn, proliferating states must seek alternative means to develop their weaponry. But since the weapons inspectors close off the most efficient means, the alternative means are inherently costlier, essentially transforming the bargaining environment to the Pareto dominant range.[11]

Consequently, imperfect monitoring is perfectly acceptable. Weapons inspectors do not act isolation; they are part of a greater negotiation strategy. While the cost of proliferation might still prove worth the finished product,

[11] Of course, this assumes that weapons inspectors are costless. (Figure 8.5 made the same assumption.) In practice, this is not the case. However, for the Theorem to still hold, we only need that inspectors cost less than the inefficiency from war and weapons construction, which seems intuitively accurate. In the Iraq case discussed later, this clearly held true: UNSC Resolution 986 funded weapons inspections with 0.8% of the revenue brought in from the food-for-oil program, a pittance compared to the enormous cost of war.

the additional inefficiency facilitates butter-for-bombs agreements, thus removing the incentive to proliferate at all. Despite the nuisance and risk of conventional intelligence leaks, rising states begrudgingly accept inspectors to make their nonproliferation commitments credible.

So, to summarize, weapons inspections need not remove imperfect information. Indeed, Theorem 8.1 applies to a game with no direct observation and no noisy signal about past moves. Rather, inspections raise proliferation costs and in turn make proliferating a dominated strategy. Declining states, aware that proliferation is not rationalizable, know that the rising state will not build. In turn, the declining state chooses not to prevent. War does not occur despite imperfect information remaining throughout.

For practical purposes, it is important to note that weapons inspectors appear completely worthless under such conditions. Given the nature of their jobs, one might expect weapons inspectors to discover actual weapons. But if inspectors were to find such violations, their hosts would simply not invite them in the first place. In turn, when unobstructed inspectors hit the ground, they fail to yield any new insights that the invitation itself did not already reveal. *Weapons inspectors never find weapons.*

As such, it is inappropriate to evaluate the usefulness of weapons inspections by pure observation. Rather, we can only appreciate their benefits by considering the counterfactual case. Although inspections merely act as a nuisance, the inconvenience is remarkably valuable because it makes violating arms agreements counterproductive. Thus, the presence of the weapons inspectors paradoxically make the weapons inspectors unnecessary.

8.5 Understanding the Origins Iraq War

Using a related model, Debs and Monteiro (2013) claim that imperfect information caused the 2003 Iraq War. This section reviews the run up to the conflict and shows that Iraq attempted to increase its cost burden to develop prohibited weapons, hoping to make its nonproliferation commitments credible and avoid fighting. This matches the logic of Theorem 8.1. The United States ignored these efforts. Thus, *imperfect information did not cause the Iraq War.*

We break up the argument into two parts, beginning with Iraq's behavior between the 1991 Gulf War and September 11 terrorist attacks and ending with the strategic decisions made before the 2003 Iraq War started.

8.5.1 Strategic Jockeying: 1991-2001

To understand the build up to the Iraq War in 2003, the end of the Persian Gulf War in 1991 serves as a good starting point. The American-led coalition had trounced Saddam's forces over three days of ground fighting; Iraqi losses were counted by the tens of thousands while coalition losses were merely counted by the hundreds. United Nations Security Council Resolution 687 laid out the terms of the peace. Iraq was to cease all nuclear weapons development, eliminate its biological and chemical weapon stockpiles, and limit its ballistic missiles to a 150km range.

History showed that the terms were not self-enforcing, and the international community would struggle with maintaining Iraqi compliance for the next decade. In terms of weapons development, Saddam maintained a status quo advantage. No matter the sanctions levied against him, Saddam could develop weapons all he wanted to–the United Nations would have to raise another coalitional force to stop him. And given Iraq's turbulent relationship with both Iran and the West in the past five years, keeping existing weapons programs going looked tempting.

Moreover, political realities at the end of the Gulf War hinted that Saddam could skirt his treaty obligations without serious consequence. Despite the blow out Iraqi troops suffered during the short ground war, the United States decided not to invade Baghdad and remove Saddam from power. This restraint resulted from two practical concerns. First, attempting to do so risked breaking up the diverse coalition. And second, although the United States enjoyed a major technological advantage in open air and open ground combat, Baghdad would be an urban front. American causalities might have spiked accordingly. So George H.W. Bush settled short of total victory. And yet, to assuredly stop Saddam from expanding his weapon stockpiles, such an invasion of Baghdad would be necessary. UNSCR 687 was a signal of American weakness–if the U.S. were truly committed to regime change, Saddam figured they would have toppled him from power in February of 1991.

As such, Saddam systematically skirted his treaty obligations. His gambit worked in the short term, as the United States responded tepidly. Iraqi agents frustrated the United Nations Special Commission on Iraq for the next seven years, lying about existing stockpiles and limiting inspector access throughout the country. The U.N. mostly responded by imposing sanctions. Finally, in December 1998, inspectors evacuated Iraq and the United States launched Operation Desert Fox, a four day aerial bombing campaign on suspected weapons facilities. Iraq responded by formally declaring that the Special Commission would not be welcomed back.

The status quo persisted for the next few years. Western intelligence presumed Saddam continued building up his military infrastructure, and United Nations remained outside of Iraq.

8.5.2 Costly Signaling in the Post-9/11 Era

The September 11 terrorist attacks refocused American priorities abroad. After a successful run in Afghanistan, President George W. Bush turned his focus to Iraq. Whereas the Clinton administration had found the quarantine tolerable, Bush quickly threatened a full-scale intervention if the Iraqi situation did not improve. The United Nations correspondingly created the United Nations Monitoring, Verification, and Inspection Commission (UNMOVIC) to handle a new round of weapons inspections. Swedish diplomat and former head of the IAEA Hans Blix came out of retirement to head the team.

But even with the international attention, Bush's threat was not inherently credible. While the previous section covered American hesitation in Iraq, Saddam viewed the events of the post-Vietnam era as evidence that the United States had a weak stomach for war in general. Indeed, Saddam was

> aware that after a few casualties in Lebanon and Somalia, the Americans had retreated with their tails between their legs....In Somalia, the United States had disengaged after suffering eighteen dead in the streets of Mogadishu. In Kosovo, the U.S. had relied on a seventy-eight-day air war and prided itself that it had won without a single casualty. (Gordon and Trainor 2006, 66)

Despite having repeatedly violated the terms of the Persian Gulf War peace treaty in the decade earlier, the American response had been consistently weak in his eyes. Why would this time be any different? Even after 9/11, Iraq was not an obvious extension in the war on terrorism. From Saddam's perspective, it was unclear whether the American public would accept a second front, especially since an invasion of Iraq would likely come at a greater cost than the invasion of Afghanistan. Meanwhile, on the diplomatic front, the United States could not muster an alliance on par with the one in

Afghanistan or the force that expelled Saddam from Kuwait in 1991. So not only would the overall cost burden be greater, but the United States would have to absorb a greater share of it.

Furthermore, Iraq had plenty to lose by revealing its lack of serious power. Under the status quo, Saddam could bluff strength against (and perhaps concessions from) Iran and quell potential domestic uprisings from the Shiite population (Lake 2011, 30-31). If Iraq opened to weapons inspections, this "deterrence by doubt" strategy would fold (Gordon and Trainor 2006, 64-65)–conceding to the Bush administration's demands would reveal Iraqi weakness, leave Iraq powerless in the region, and perhaps encourage domestic foes. It would also prevent Iraq from developing serious programs to deter through real force later rather continuing with the bluff. So Saddam resisted to the U.N.'s initial calls for a new round of weapons inspections. And, as late as January 2003, Iraq appeared aloof about the gravity of the situation (Gordon and Trainor 2006, 132).

By this point, in terms of this chapter's model, the situation is a good match for the parameters of Proposition 8.4. Iraq found weapons of mass destruction potentially worth the investment cost, as that strength would have partially insulated him from external coup attempts and raised his status relative to Iran. In addition, Iraq knew the United States viewed preventive war as preferable to an Iraqi buildup. Meanwhile, the U.S. was unwilling to offer inducements to convince Iraq to forgo development–at this point, Iraq could easily evade inspections and proliferate in relative peace. Thus, mutually preferable butter-for-bombs agreements were unavailable.

However, by February 2003, it became clear to Saddam that the United States would invade if he failed to comply. The U.S. had stationed 100,000 troops in Kuwait. This was not the meek signal the Clinton administration sent in 1998. Rather, it was the type of costly signal that credibly revealed Bush's resolve to invade in the aftermath of 9/11. And as standard signaling theory would predict, Saddam internalized the threat and backed down.[12] Deterring Iran was important, but not as important as staying in power–and, by extension, staying alive.

In the absence of endogenous institutions, accidental preventive war is rational under the pre-February 2003 conditions. But realizing that war could be on horizon, Saddam made a push for credibility as Theorem 8.1 would predict. Iraq's relationship with weapons inspectors changed accordingly. From

[12]See Fearon 1994, Fearon 1997, Slantchev 2005, and Slantchev 2011.

the inception of UNMOVIC, Iraqi duplicity had frustrated Blix's inspectors just as they had frustrated inspectors from the previous commission. But that changed by March 2003. Blix noted the remarkable turnaround:

> Iraqi counterparts had toward the end become almost frantic in submitting material, seeking and finding persons [they] could interview....Had there been denials of access? Any cat-and-mouse play? No. Had the inspections been going well? Yes. (Blix 2004, 10-11)

Moreover, UNMOVIC interrogators could now interview persons of interest without any recording devices or Iraqi officials present (Blix 2004, 209). Even Saddam's presidential palaces–previously the most protected locations–became open for inspection (Blix 2004, 14).[13]

Was this inconvenience enough to make Saddam's commitment to nonproliferation credible? As long as weapons inspectors remained on the ground–and UNMOVIC was prepared to stay for the long-term (Blix 2004, 11)–Iraq's path to weapons of mass destruction would have been long, arduous, and (most importantly) costly. Given Saddam's fragile bargaining position, this was likely enough to disincentivize proliferation.

However, to be careful not to select on the dependent variable, it is important to substantively interpret the cutpoints under which the theory applies. Recall that Proposition 8.3 shows that butter-for-bombs agreements are credible if $p'_R - p_R < c_D + c_R + \frac{k(1-\delta)}{\delta}$. Reworking this in terms of the cost of proliferation, $k > \frac{\delta(p'_R - p_R - c_D - c_R)}{1-\delta}$ must hold. Thus, as long as $p'_R - p_R - c_D - c_R$ is sufficiently small, butter-for-bombs succeeds.

Now consider these terms in isolation. In the Iraq War case, the extent of the power shift $(p'_R - p_R)$ was reasonable; Saddam was virtually powerless

[13]Debs and Monteiro (20) claim that Blix was "remarkably ambiguous" in his public statements just prior to the war. This is because Blix viewed UNMOVIC's role as a fact-finding mission. He was thus unwilling to make declarations on whether Iraq had complied with the "immediate" cooperation demand the Council had issued (Blix 2004, 210). Rather, he felt the onus was on Security Council members to consider his findings and then rule whether Iraq had abided by UNSC Resolution 1441. This distinction is clear in his final presentation to the Council, in which he stated that Iraq's efforts post-February 2003 could "be seen as active, or even proactive," but "these initiatives three to four months into the new resolution cannot be said to constitute 'immediate' cooperation" as required by Resolution 1441 (Blix 2003). As such, Blix was *unambiguous* in his belief (and in his presentation) that Iraq was cooperating after the Bush administration had sent its costly signal in February.

at the time (leaving p_R close to 0) but could have protected himself against Iran in a post-shift stage and warded off any threat of regime change.[14] As such, p'_R would equal a positive, albeit still small, amount.

Next, the United States' costs of war (c_D) has two interpretations. The first is that the Bush administration knew that the cost of combating a insurgency would be high but ignored them for domestic political purposes, perhaps to continue the image of Bush as a strong wartime president heading into the 2004 election. If this is the case, then the explanation for the Iraq War breaks from the rational unitary actor approach, and our work is done–imperfect information did not cause the war.

The second interpretation gives the Bush administration the benefit of the doubt. Perhaps the Bush administration severely underestimated the costs of war.[15] If so, for the purposes of the contemporary bargaining problem, conventional wisdom believed the United States "would, given the precedent of the 1991 Gulf War, no doubt prevail" and that "the U.S. ability to depose Saddam promptly was never in doubt" (Debs and Monteiro 2013, 17).[16] Given the perceived low costs of fighting, peace seems precarious.

Yet even under that assumption, there is still reason to question whether imperfect information caused the war–the same forces that ensured small American costs also guaranteed catastrophic costs for Saddam Hussein. Given Saddam's despotic past, compliance problems, and the Bush administration's view that democracy could spread throughout the Middle East, there could be little doubt that the United States would pursue a policy of absolute war and regime change.[17] For Saddam, this meant that war implied life imprisonment at best and death at worst.[18] Consequently, we would imagine the value for c_R as quite large, likely large enough that Saddam would have preferred

[14]In addition, other forms of WMD could have given Saddam minor gains against domestic threats.

[15]Coe (2012) argues that the United States' cost for war was actually negative.

[16]For a while, Saddam held out hope that international pressure from Russia and France combined with Iraq's ability to mount *some* costs against American troops would convince America to back off. By the eve of war, however, it was clear that the United States did not care enough about their Security Council peers' judgment to deter conflict (Lake 2011, 32).

[17]Wolford, Reiter, and Carrubba (2011) explore a model with incomplete (rather than imperfect) information and post-treaty compliance problems.

[18]Indeed, autocratic leaders are much more likely to face post-war punishment (Debs and Goemans 2010; Goemans 2000; Chiozza and Goemans 2011).

giving the United States the entirety of the bargaining good than fight.[19] That drives down the value of the denominator, making the cost threshold $k > \frac{\delta(p'_R - p_R - c_D - c_R)}{1-\delta}$ easier to fulfill.[20]

So, on the eve of war, the United States had the opportunity to buy off Saddam. Because weapons inspectors are not perfect, Saddam could have taken those concessions and kept proliferating in secret. However, inspections greatly limited Saddam's options:

> The bulk of [IAEA] inspections were at state-run or private industrial facilities, research centers, and universities–focusing on locations where we knew Iraq had maintained significant technical capabilities in the past, or on new locations suggested by the analysis of open-source information, or on facilities that were identified through satellite imagery as having been modified or constructed since 1998. (ElBaradei 2011, 59)

In other words, inspectors had flooded all of the desirable proliferation locations. Given the extra cost required to proliferate, if the United States made the properly-sized bribe, building would have been a dominated action. Thus, imperfect information or not, the United States could trust that Saddam would not proliferate–not for a lack of desire but rather because his incentives no longer gave him reason to do so.[21]

All told, it appears that the inconvenience of fully functional weapons inspectors was sufficient to make Saddam's nonproliferation commitment credible, so long as the United States provided the correct inducements as well.

[19]Further, Saddam's concerns of CIA infiltration of weapons inspections became unimportant. A slower death through intelligence leaks or assassination attempts was preferable to immediate regime change.

[20]Tarar (2013) explores a model in which military mobilization can cause a commitment problem and lead to war. However, his logic appears not to apply to the February 2003 mobilization–war only occurs if the costs of war are smaller than the power shift caused by the mobilization. Yet, for the reasons outlined above, Saddam's costs for war likely left open a bargaining range.

[21]Moreover, this chapter has assumed that weapons inspections provide absolutely no informational context but rather inflate costs of proliferation. If inspections served both purposes, stronger signals would again push the actors into the Pareto improving parameter space of Proposition 8.3. This is because a revealed weapons program leads to preventive war. In turn, any decision to proliferate carries with it an implicit cost of the wasted cost with positive probability. As a result, the rising state is willing to accept smaller butter-for-bombs offers as the strength of the signal increases.

If "known unknowns" were the cause of potential conflict between the United States and Iraq, we would then expect war to have occurred pre-February 2003 and not afterward. Yet the opposite is true–the United States maintained the peace while the situation in Iraq was most opaque. When Saddam finally opened the country to useful inspections, the Bush administration ignored the change and pressed forward.[22]

As a final note, this chapter does *not* claim the Iraq War cannot be explained through rationalist approach. It is neutral in this regard. Although the U.S. intelligence agencies failed to appreciate Iraqi efforts to disarm, those efforts are irrelevant if a different bargaining problem caused the invasion. Further, the Bush administration's push to the United Nations could have been to rally domestic support for the war (Chapman and Reiter 2004; Chapman 2011; Grieco et. al. 2011) or secure assistance from allies.

Instead, the goal for this section was to argue a negative, namely that the cause of the Iraq War was not imperfect information. Consequently, more research needs to be done on the cause of the Iraq War–whether the focus is on rationalist unitary actor explanations, rational domestic politics explanations, or behavioralist explanations.[23]

8.6 Conclusion

The nonproliferation regime faces a seemingly insurmountable task. Nuclear weapons afford their owners unparalleled coercive power. Worse, covert proliferation programs allow states to progress undetected, preventing competing states from taking appropriate countermeasures. The incentives appear ripe for proliferation and its associated inefficiency.

However, good institutions morph perverse incentives to allow for cooperative behavior. If proliferation comes too easy, declining states cannot buy

[22]To wit, in preparation for inspections, Saddam cleaned older WMD sites. American satellites, monitoring Iraqi movements, recognized the effort. But rather than hailing this as a victory–the sites had been neutered–Colin Powell's presentation to the Security Council on February 5 cited it as a coverup (Gordon and Trainor 2006, 134).

[23]Lake's (2011) call for a behavioralist focuses on how suboptimal strategizing from leaders can cause conflict. However, the narrative from this section suggests an alternative behavioralist approach. If the Bush administration pursued war for domestic political purposes, the leadership nevertheless gained support from the public by appealing to rationalist unitary actor explanations. Behaviorally, then, we should better understand how the public fails to properly assess international unitary actor negotiations.

off potential rising states in a cost-effective manner; instead rely on preventive war to stifle possible power shifts. But because such preventive war sometimes strikes innocent states, those potential rising powers have incentive to increase the cost necessary to proliferate. This lowers the necessary bribe declining states must give, which in turn make the commitment not to proliferate credible. Both states enjoy the additional surplus created from peace and lower levels of weapons investment.

While this chapter has shown that weapons inspectors build trust and lead to cooperation, we still have cases where proliferating states forbid inspections and flagrantly violate nonproliferation commitments. Given the benefits of butter-for-bombs treaties and institutionalizing compliance to make such agreements credible, why does this occur? The next chapter introduces incomplete information to the interaction and shows that declining states rationally gamble and act tough, hoping that the threat of preventive war is sufficient to induce compliance. In turn, the model will help explain the critical problems the weapons inspectors encounter that this chapter's model could not adequately address.

8.7 Appendix

We now cover the proofs missing from the main text of this chapter. As stated earlier, the proof for Proposition 8.1 follows from the proof for Proposition 3.2, so we omit it here. Thus, throughout this appendix, we only consider cases in which $p'_R - p_R > \frac{k(1-\delta)}{\delta}$.

8.7.1 Proof for Proposition 8.2 and Proposition 8.3

We presented Proposition 8.2 and Proposition 8.3 separate for expositional purposes. The proof, however, is identical. From Figure 8.4, note that if $x \geq p'_R - c_R - \frac{k(1-\delta)}{\delta}$, not building dominates building for R. By iterated elimination of dominated strategies, D does not prevent if $x < p_R + c_D$.[24] R

[24] If $x > p_R - c_R - \frac{k(1-\delta)}{\delta}$, iterated elimination of strictly yields a unique solution. If there is only equality, we assume that R accepts with probability 1 without loss of generality. We may do this here for the same reasons that we may do this for the standard ultimatum game, which has no equilibria in which the receiver rejects with positive probability when indifferent.

receives x and D receives the remainder. Since D's payoff is strictly decreasing in x, its optimal offer in this range is $x = p'_R - c_R - \frac{k(1-\delta)}{\delta}$.

If $x > p_R + c_D$, D prevents. If $x = p_R + c_D$, D is indifferent between preventing and not preventing. In either case, D's payoff is less than if it offered the optimal butter-for-bombs amount.

Now consider any offer smaller than $p'_R - c_R - \frac{k(1-\delta)}{\delta}$. This can only be more profitable for D if D does not prevent and R does not build with positive probability. (If D prevents, it receives its war payoff; if D does not prevent and R builds, D receives more from having induced R to accept the optimal butter-for-bombs offer.) Such an outcome cannot occur with certainty, as R could profitably deviate to building. It also cannot be the case that D does not prevent with certainty and R mixes between building and not building, as R's indifference condition requires $x = p'_R - c_R - \frac{k(1-\delta)}{\delta}$. The only remaining possibility is where D mixes between preventing and not preventing. However, the indifference conditions means that D must earn $1 - p_R - c_D$ in expectation, but this is worse than the optimal butter-for-bombs offer.

Thus, D optimally offers $x = p'_R - c_R - \frac{k(1-\delta)}{\delta}$. □

8.7.2 Proof for Proposition 8.4

Proposition 8.4 stated that D receives its war payoff if $p'_R - p_R > c_D + c_R + \frac{k(1-\delta)}{\delta}$. However, this can occur in many different ways depending on the offer D makes initially. Since there many different possibilities, we break up the cases into the three lemmas below.

Lemma 8.1. *If $x < p_R + c_D$, D prevents with probability $\frac{\delta x - \delta(p'_R - c_R) + k(1-\delta)}{\delta x - \delta(p'_R - c_R)}$ and R builds with probability $\frac{x - p_R - c_D}{\delta(x - p'_R + c_R)}$.*

Proof: If $x < p_R + c_D$, it is trivial to show that no pure strategy Nash equilibria exist and that the only way one player can be willing to mix is if the other also mixes. Thus, we look for mixed strategy Nash equilibria by deriving the indifference conditions.

To begin, let σ_p be the probability D prevents. Then R's expected utility for building equals:

$$\sigma_p[p_R - c_R - k(1-\delta)] + (1-\sigma_p)[(1-\delta)x + \delta(p'_R - c_R) - (1-\delta)k]$$

$\sigma_p[p - c_R - (1 - \delta)x - \delta(p'_R - c_R)] + (1 - \delta)x + \delta(p'_R - c_R) - (1 - \delta)k$

And R's expected utility for not building equals:

$$\sigma_p(p_R - c_R) + (1 - \sigma_p)x$$
$$\sigma_p(p_R - c_R - x) + x$$

Setting these two equations equal to each other and solving for σ_p yields:

$$\sigma_p = \frac{\delta x - \delta(p'_R - c_R) + (1 - \delta)k}{\delta x - \delta(p'_R - c_R)}$$

Note that because $x < p_R + c_D < p'_R - c_R - \frac{k(1-\delta)}{\delta}$, the numerator is less than zero. And because the denominator is less than the numerator, the denominator is negative as well. Since the denominator's magnitude is greater than the numerator's, the probability distribution is valid.

Now consider D's indifference condition. D's expected utility for preventing is $1 - p_R - c_D$ regardless of R's strategy. Meanwhile, letting σ_b be the probability R builds, D's expected utility for not preventing equals:

$$\sigma_b[(1 - \delta)(1 - x) + \delta(1 - p'_R + c_R)] + (1 - \sigma_b)(1 - x)$$
$$\sigma_b[\delta(x - p'_R + c_R)] + 1 - x$$

Setting these two values equal to each other and solving for σ_b yields:

$$\sigma_b = \frac{x - p_R - c_D}{\delta(x - p'_R + c_R)}$$

Since $x < p_R + c_D < p'_R - c_R - \frac{k(1-\delta)}{\delta}$, both the numerator and the denominator are negative. In addition, the magnitude of the denominator is greater than the magnitude of the numerator by assumptions of the parameter space in Chapter 3.

Note that if D is mixing, it is indifferent between preventing and not preventing. Since preventing yields a flat $1 - p_R - c_D$, D receives its war payoff in this case. □

Lemma 8.2. *If $x = p_R + c_D$, D prevents and R does not build.*

Proof: Suppose $x > p_R + c_D$. We can prove Lemma 8.2 through iterated elimination of strictly dominated strategies. D strictly prefers preventing if R does not build. If R builds, D earns $1 - p_R - c_D$ for preventing and

$(1-\delta)(1-x)+\delta(1-p'_R+c_R)$. Using $x = p_R + c_D$ as an upper bound for D's payoff, preventing is strictly better if:

$$1 - p_R - c_D > (1-\delta)(1-p_R-c_D) + \delta(1-p'_R+c_R)$$
$$p'_R - p_R > c_D + c_R$$

Recall that $p'_R - p_R > c_D + c_R + \frac{k(1-\delta)}{\delta}$ for this parameter space. So the inequality holds. Therefore, preventing is always strictly better than not preventing for D.

By iterated elimination of strictly dominated strategies, R does not build. This leads to a pure strategy outcome in which D receives $1 - p_R - c_D$. $\square$

Lemma 8.3. *If $x > p_R + c_D$, R does not build and D prevents with probability $\sigma_p \in [\frac{\delta x - \delta(p'_R - c_R) + k(1-\delta)}{\delta x - \delta(p'_R - c_R)}, 1]$.*

Preventing now weakly dominates not preventing. Not building is the best response to preventing, so preventing and not building is an equilibrium.

No other pure strategy Nash equilibria exist. Not preventing and not building is not an equilibrium; R needs at least $p'_R - c_R - \frac{k(1-\delta)}{\delta}$ to not want to deviate to building, but $p_R + c_D$ is less than that. Building and not preventing is not an equilibrium, as D prefers to prevent. Lastly, preventing and building is not an equilibrium, as R could deviate to not building and save the investment cost.

Now consider mixed strategy Nash equilibria. Since preventing weakly dominates not preventing, if R mixes, D must prevent as a pure strategy. But since R strictly prefers not building in this case, R cannot optimally mix.

Thus, the remaining cases to consider require D to mix and R to select a pure strategy. D is indifferent between preventing and not preventing only if R does not build. For R to be willing to not build, Lemma 8.1 shows that $\sigma_p \geq \frac{\delta x - \delta(p'_R - c_R) + k(1-\delta)}{\delta x - \delta(p'_R - c_R)}$.

Thus, in equilibrium, R does not build and D prevents with probability at least $\frac{\delta x - \delta(p'_R - c_R) + k(1-\delta)}{\delta x - \delta(p'_R - c_R)}$. Since D is indifferent between preventing and not preventing in all of these cases, it earns its war payoff. $\square$

To conclude the proof of Proposition 8.4, the final step is to find the value for x that maximizes D's payoff in the subgames the lemmas described. But note that D receives $1 - p_R - c_D$ regardless of its choice. As such, *all* offer sizes are optimal. In turn, a SPE supports any value for x, and SPE exist in which proliferation occurs, preventive war occurs, and accidental preventive war occurs. □

8.7.3 Proof for Theorem 8.1

Recall from Proposition 8.4 that when $p'_R - p_R > c_D + c_R + \frac{k(1-\delta)}{\delta}$, or $k < \frac{\delta(p'_R - p_R - c_D - c_R)}{1-\delta}$, multiple equilibria exist. The proof for Proposition 8.4 showed that the Pareto efficient equilibrium from that set was still inefficient. Namely, in that equilibrium, D mixes between preventing and not preventing and R mixes between building and not building. The indifference condition for D implies that D receives its pre-shift war payoff, or $1 - p_R - c_D$, for the game. Since the sum value of the game is at most 1, this means that R receives no more than $p_R + c_D$. But because mixing entails positive probability of both war and investment (and the corresponding deadweight loss of c_D, c_R, and k), it must be that R receives strictly less than $p_R + c_D$.

The size of that deadweight loss depends on the offer D makes initially, and there are infinitely many such possibilities in equilibrium. Nevertheless, the exact portion is irrelevant to the argument. We can instead write the amount of deadweight loss as $L > 0$. Consider equilibrium conditions for $\hat{k} > \frac{\delta(p'_R - p_R - c_D - c_R)}{1-\delta}$. Propositions 8.2 and 8.3 state that D offers $x = p'_R - c_R - \frac{\hat{k}(1-\delta)}{\delta}$ and R accepts. D prefers this outcome if:

$$1 - p'_R + c_R + \frac{\hat{k}(1-\delta)}{\delta} > 1 - p_R - c_D$$

$$\hat{k} > \frac{\delta(p'_R - p_R - c_D - c_R)}{1-\delta}$$

Meanwhile, R prefers the alternative outcome if:

$$p'_R - c_R - \frac{\hat{k}(1-\delta)}{\delta} > p_D + c_D - L$$

$$\hat{k} < \frac{\delta(p'_R - p_R - c_D - c_R + L)}{1-\delta}$$

Such a mutually preferable alternative investment cost $\hat{k}$ exists if:

$$\frac{\delta(p'_R - p_R - c_D - c_R)}{1-\delta} < \frac{\delta(p'_R - p_R - c_D - c_R + L)}{1-\delta}$$

$$L > 0$$

This holds, since the deadweight loss is a strictly positive amount. Therefore, both parties benefit from increasing the costs of proliferation. $\square$