

# Why Give Nuclear Assistance to Stop Would-Be Proliferators?

William Spaniel\*

July 10, 2018

## Abstract

Members of the nonproliferation regime give technical assistance to countries contemplating nuclear weapons. This is puzzling: it facilitates the behavior donors wish to stop, and other forms of concessions do not have this drawback. Why do it? I develop a model of uncertainty, bargaining, and nuclear proliferation. In it, assistance hastens acquisition time but also generates a signal about the recipient's domestic nuclear capacity. This allows donor states to better calibrate other concessions to the recipient. In equilibrium, donor states sometimes find the information worth sacrificing bargaining leverage. However, despite providing information, assistance can cause proliferation if donors believe the recipient is competent but observe a false signal indicating incompetence. Strategic assistance from international institutions does not suffer from this problem, though a selection effect can make observational data appear to suggest otherwise.

---

\*Department of Political Science, University of Pittsburgh, Posvar Hall 4446, Pittsburgh, PA 15260. (williamspaniel@gmail.com, <http://williamspaniel.com>).

# 1 Introduction

On December 8, 1953, still during the dawn of the nuclear era, President Dwight D. Eisenhower outlined the foundations of “Atoms for Peace” in front of the United Nations General Assembly. His premise was simple. To disincentivize further weapons proliferation, capable states ought to offer peaceful nuclear technology to compliant states. Both the International Atomic Energy Agency (IAEA) and the Nuclear Non-Proliferation Treaty trace back to the speech, and it remains a central tenant of the nonproliferation regime today.

Nevertheless, Atoms for Peace is puzzling. To convince other countries not to proliferate, donor states make nuclear proliferation cheaper and easier. That powerful countries wish to give concessions to others to ensure nonproliferation is not surprising and is a key component of nonproliferation theory. What is surprising is that they choose to use nuclear technology as a carrot instead of other policy or military concessions. Obtaining compliance in coercive bargaining requires giving opponents concessions commensurate with their outside options. By enhancing a country’s nuclear proficiency, donor states improve recipient states’ ability to proliferate. This would seem to force donors to give further concessions down the road or result in more countries proliferating outright. Indeed, it seems that all Atoms for Peace does is “spread temptation” (Fuhrmann, 2009*a*).

Why, then, do donor states transfer nuclear technology? I argue that those specific concessions help resolve an information problem that donors face. To convince would-be proliferators to end nuclear weapons programs, concerned states need to offer a deal. However, due to the technical nature of nuclear proliferation, intelligence estimates of foreign countries’ costs and development speeds of nuclear weapons are noisy (Montgomery and Mount, 2014). This hinders a donor’s ability to offer the right quantity of concessions to convince would-be proliferators to forgo nuclear weapons. Fortunately, integrating nuclear programs allows donor states to better understand a recipient’s domestic proficiency. Learning this information is sometimes worth the drawback of facilitating further proliferation progress, thus explaining the Atoms for Peace strategy.

To develop this argument, I build a bargaining model of nuclear proliferation, aid, and uncertainty. A donor state does not know a recipient state’s underlying nuclear proficiency. It begins by choosing whether to transfer nuclear technology to the recipient

state. Doing so enhances the recipient state’s proficiency but also gives the donor state an informative—but noisy—signal about the recipient’s natural skill. Afterward, the donor chooses to offer policy concessions. The recipient either accepts or rejects the offer. Rejecting leads to proliferation, which comes cheaper and more quickly if the recipient received aid earlier.

The formal analysis indicates that the aforementioned intuition is correct: donor states provide aid to increase their information. However, the model also reveals some counterintuitive implications. Summarizing the conflict literature, Kydd (2010) writes that “[i]f uncertainty leads to cooperation failure, then information can lead to conflict resolution.” Yet increasing information here sometimes lead to more proliferation, not less. Why? If aid were impossible, the donor’s initial beliefs might suggest that the recipient is a proficient type, leading to generous offers that would be accepted by all types. Therefore, no proliferation would occur under those circumstances. In contrast, when making a transfer, the donor may sometimes receive a signal indicating incompetence. This can cause the donor to downgrade its offer and lead to proliferation if the signal was misleading.

I then take the model in a different direction. Nuclear assistance extends beyond the realm of bilateral accords. International institutions, and the IAEA’s Technical Cooperation initiative in particular, also make transfers. These institutions have a different set of incentives. They do not make coercive bargaining concessions to would-be proliferators and instead just wish to minimize proliferation behaviors. I therefore retool the model to investigate how an institution’s transfer strategy differs from a rival state’s strategy.

Three key results emerge. First, when the opponent of the recipient is initially worried about the latter’s competence, the institution has *less* incentive to provide assistance. This is because the opponent would naturally make safe offers that induce nonproliferation, and any information the institution might generate can only interfere with that. Second, when the opponent of the recipient is initially skeptical of the latter’s competence, the institution has *more* incentive to provide assistance. This is because proliferation would often result in the absence of additional information, and the institution does not care that assistance forces the opponent to increase its offers. Finally, aside from the signal, the institution may wish to increase the recipient’s skill to make both competent and incompetent types behave more similarly. This has a

second-order effect of mitigating the opponent's information problem and can result in less proliferation even if the signal is not informative.

The model also generates important empirical implications for scholars and policymakers alike. Prior research shows a correlation between atomic assistance and proliferation behaviors. A reasonable conclusion is that transfers *cause* more proliferation. The natural policy implication is to reduce or restrict assistance programs (Fuhrmann, 2009*a*, 39-41; Brown and Kaplow, 2014, 421-422). My results suggest an alternative interpretation: there may be a selection problem at work. Assistance can occur when proliferation is most likely in its absence. Despite transfers reducing the barriers to development, a blanket ban on assistance could counterintuitively result in *more* proliferation.

The next section connects this paper to the broader nuclear proliferation literature. However, from a formal perspective, it is closest to Arena and Wolford (2012). Their paper explores investment in military intelligence in the shadow of crisis bargaining. They find that some level of investment may be optimal and that learning may increase the probability of war despite reducing information problems on average. My work differs on three dimensions. First, the substantive issues are completely different. Second, my intelligence gathering mechanism improves the *opponent's* outside option, making it unclear whether the donor would still want to pay for the information. Finally, I also investigate how an international organization's assistance decision may differ from a state involved in coercive bargaining.

From a broader perspective, my results also contribute to the literature on international institutions. International relations theorists commonly argue that a main function of institutions is to provide information (e.g., Keohane, 1984). This information then facilitates negotiations between states, allowing them to reach bargains they would otherwise not agree to. My results suggest that such a mechanism may not generalize. In fact, the institution in my model sometimes fares *worse* by providing information, even if information acquisition is costless.

## 2 Microfoundations: Bargaining, Assistance, and Signals

Nuclear assistance is common across the world, though the scope and donor processes vary from project to project. Two types receive the most attention. First, Nuclear Cooperation Agreements (NCAs) are bilateral deals between two states to share nuclear technology, knowledge, or infrastructure. Between 1950 and 2000, states signed 2,470 NCAs for peaceful nuclear applications, with a slight trend upward over time (Fuhrmann, 2009*b*, 193-194). Meanwhile, one of the International Atomic Energy Agency’s central functions is to integrate the nuclear science community and facilitate the exchange of information and training of personnel.<sup>1</sup> IAEA technical cooperation projects, in which IAEA scientists assist recipient states, are far more common than NCAs, with 890 ongoing in 2010 alone (Brown and Kaplow, 2014, 405).

These projects do not intend to assist countries in developing nuclear weapons.<sup>2</sup> However, nuclear technology for civilian applications often has dual uses. For example, assisting a country with their nuclear power is ostensibly benign. But work in that area can train scientists in basic nuclear principles that they could eventually use in the construction of a weapon. As a result, assistance allows countries to build nuclear weapons faster and for a lower price.

The quantitative “supply side” literature indicates that this is a recipe for more proliferation. Based on simple cost-benefit theories, countries with cheaper and faster routes to nuclear weapons are more likely to proliferate. After all, the gains come quicker and the investment is more manageable. The data reflect this. Countries with higher industrial infrastructure are more likely to pursue and acquire nuclear weapons (Singh and Way, 2004). So too are countries with higher nuclear proficiency levels (Jo and Gartzke, 2007; Smith and Spaniel, 2018), which transfers conceivably increase. More directly, countries receiving more nuclear cooperation agreements and technical cooperation assistance exhibit the same patterns (Fuhrmann, 2009*a*).

These theoretical and empirical results suggest an obvious question, and one that (Fuhrmann, 2009*a*, 41) concludes with after establishing the connection: “[w]hy do sup-

---

<sup>1</sup>See Article III A 3-4 and Article VIII: <https://www.iaea.org/about/statute>

<sup>2</sup>This is the realm of “sensitive” nuclear assistance, which also predicts high rates of proliferation (Kroenig, 2009).

pliers provide civilian assistance” given the perverse incentive? A second literature hints at a path toward an explanation. Just because a state can develop nuclear weapons does not mean it will. The nuclear path states select depends on the availability of non-proliferation inducements (Reiss, 1988; Paul, 2000; Miller, 2014; Spaniel, 2015; Bas and Coe, 2016; Debs and Monteiro, 2016; Volpe, 2017). Assistance or not, nuclear weapons are expensive.<sup>3</sup> Some states nevertheless find nuclear weapons attractive because the benefits outweigh the costs. However, at the end of the day, those benefits have an upper cap. Meanwhile, other states suffer externalities with the arrival of a new member of the nuclear club. For enemies, this includes the direct security and coercive bargaining losses. For allies, a new nuclear friend may deviate from the patron’s preferred foreign policy. Both types of states suffer the damage to the nonproliferation regime and the risks of a catastrophic accidental nuclear war. Correspondingly, opponents often offer tacit or explicit nonproliferation agreements, trading concessions today for the termination of a would-be proliferator’s nuclear program.

However, the existence of mutually beneficial agreements does not guarantee their implementation (Fearon, 1995). Particularly problematic for nuclear negotiations is uncertainty over a state’s nuclear proficiency. The United States, despite pouring substantial sums of money into intelligence programs, has a poor track record in predicting other states’ development speeds (Montgomery and Mount, 2014). This is because most nuclear progress depends on solving technical problems. The ability to accomplish these tasks lies within the brain trust of a state’s nuclear industry. Without intimate knowledge of key nuclear scientists’ skill, outsiders can only make crude nuclear timetable estimates.

This uncertainty creates second-order bargaining problems. Nonproliferators have competing incentives. On one hand, they want to offer sufficient concessions to induce compliance and avoid suffering the externality. But conditional on securing compliance, they want to offer the smallest deal possible to avoid sacrificing unnecessary concessions. The deal must be commensurate with the speed and cost of nuclear weapons. Thus, if the nonproliferator is uncertain of its opponent’s proficiency, it may offer a deal insufficient to induce acceptance. Worse, simple communication is an ineffective means

---

<sup>3</sup>Schwartz (2011) estimates the United States’ total nuclear expenditures on construction, deployment, and command and control at almost \$4.5 trillion between 1940 and 1996, in constant 1996 dollars.

to resolve the issue. After all, if competent states require more concessions, incompetent states have an incentive to misrepresent their skill level.

Calculation of this minimum offer indicates another problem with providing nuclear assistance. Proficiency hastens development times and reduces the cost of proliferation. Thus, providing assistance would seem to sabotage the nonproliferator's bargaining position. After all, regardless of the state's native skill, transfers would force the nonproliferator to provide greater concessions to appease the would-be nuclear state.

The negotiations literature provides a potential explanation for this observation: nuclear assistance programs *are* concessions. This may have some merit. As Taiwan considered nuclear weapons, the United States threatened to withdraw technical cooperation (Miller, 2014, 931). Taiwan values these critical benefits. Nuclear power plants provide 20% of the island's power, but the United States provides the fuel (Mitchell, 2004, 305). The United States also viewed construction of the Kori-2 nuclear power plant as a concession to South Korea, one that could be revoked should Seoul not commit to nonproliferation norms (Miller, 2014, 934). Cases like these involve a donor giving the recipient a subsidy on nuclear technology. But assistance-as-concessions overlooks how donors could offer non-nuclear concessions that do not perversely lower the barriers to proliferation.

In fact, such non-nuclear concessions are common, even in cases where the would-be proliferator also received technical assistance. For example, beyond nuclear power assistance, the United States also gave \$1.5 billion in military aid in exchange for South Korea forgoing nuclear weapons (Drezner, 1999, 255). Over a twenty year period, Western countries supplied Pakistan with a research reactor, power plant, fuel, heavy water, a fuel production facility, and a heavy water production facility (Fuhrmann, 2009a, 20). In the following years, the United States attempted to convince Pakistan to forgo nuclear weapons with large offers of military aid.<sup>4</sup> Unlike South Korea, though, these efforts failed, and Pakistan proliferated.

As hinted at earlier, assistance also has an unappreciated secondary effect: it integrates the donor's nuclear field with the recipient's. Consider, for example, the IAEA's Radiation Protection Advisory Team (RAPAT). RAPAT formed in the 1980s to work with member states on safety protocols in nuclear facilities. Following its launch, "[t]he

---

<sup>4</sup>See the National Security Archive's Electronic Briefing Book # 333 (<https://nsarchive2.gwu.edu/nukevault/ebb333/>).

IAEA was surprised to learn that among its member states visited...nearly half lacked the minimum radiation safety infrastructure, more than 50 countries” (Gonzalez, 2001, 137). It then created the Model Project in Radiation Protection through the technical cooperation program to provide solutions to troubled states.

A primary consequence of these activities is lower costs and faster speeds of nuclear projects, including weapons if the recipient so desires. Indirectly, they also inform other states about a would-be proliferator’s nuclear competency. This was, in fact, a major goal of President Eisenhower’s “Atoms for Peace” program. Nuclear sharing became a source of “scientific intelligence gathering” where scientists could “assess the quality of what [other scientists] were doing” (Kirge, 2006, 166). In the RAPAT case, for instance, observers would downgrade their beliefs about a country’s skill if it does not know it should have basic safety measures in place. More broadly, interested actors can seek out information about the projects from IAEA officials or read annual reports detailing the IAEA’s assistance.<sup>5</sup> In bilateral settings, nuclear bureaucrats can learn directly from each other. For example, despite potential animosity, the Argentinian and Brazilian states agencies began working on joint projects in the 1980s. They also provided assistance to each other using their relative strengths, with Argentina supplying uranium concentrate and Brazil producing pressure vessels (Doyle, 2008). Meanwhile, the United States recognized that Taiwan had begun making progress toward a weapon by tracking its import of nuclear technology from Western states (Hersman and Peters, 2006, 544).

Wrapping up, I incorporate all of these incentives—bargaining, technical assistance, and signals—into my model. The results indicate that states and institutions sometimes find information acquisition more valuable than giving a would-be proliferator a faster and cheaper route to nuclear weapons. In comparison to existing explanations, mine has a couple of desirable features. First, Fuhrmann (2009*a*) and Brown and Kaplow (2014) recognize the perverse incentives of assistance but do not explain why actors ostensibly seeking to minimize proliferation take actions that promote it. The utility gain of information resolves that puzzle. Second, Hymans (2012) argues that assistance programs can lead to brain drain, ultimately *decreasing* a state’s competency. This makes sense but does not match the observed statistical connection between assistance and

---

<sup>5</sup>For example, [https://www.iaea.org/About/Policy/GC/GC61/GC61InfDocuments/English/gc61inf-7\\_en.pdf](https://www.iaea.org/About/Policy/GC/GC61/GC61InfDocuments/English/gc61inf-7_en.pdf).

proliferation activities. In contrast, my mechanism can counterintuitively predict more proliferation following assistance despite the information gain.<sup>6</sup>

### 3 The Model

The game consists of two states, a **Donor** and **Recipient**, negotiating over the latter’s decision to build nuclear weapons. Nature begins by drawing R as an “incompetent” type with probability  $p$  and a “competent” type with probability  $1 - p$ . R observes the realization but D does not.

D then chooses whether to provide aid or not. If it does, Nature sends D a noisy signal about R’s type, reporting it as either having “low” or “high” competence. In particular, Nature chooses the message that corresponds with the true type with probability  $.5 + q$ , where  $q \in (0, .5]$ , and chooses the misleading message with probability  $.5 - q$ . Higher values of  $q$  indicate a more informative signal.

Regardless of D’s aid decision and the message possibly received, a round of bargaining follows. Here, D offers an ultimatum  $x \geq 0$ . R can accept or reject. The game ends either way.

Payoffs are as follows. Regardless of how negotiations end, D pays  $k > 0$  if it provides assistance.<sup>7</sup> Beyond that, if R accepts the ultimatum, it earns  $x$  and D receives  $-x$ , reflecting the transfer. This is the case regardless of R’s type or any aid it may have received. If R rejects, however, both of those points come into play. Underlying skill and aid transfers affect the speed of proliferation and the cost to R. Let  $\delta(\bullet) \rightarrow [0, 1]$  map an underlying skill and transfer (or lack thereof) to a discount rate. I assume that the function is strictly increasing and that skill and aid have diminishing marginal returns.<sup>8</sup> Denote the possible transfer as  $t > 0$ , the incompetent type’s skill as  $s > 0$ ,

---

<sup>6</sup>A third explanation for NCAs is that disinterested states provide nuclear technology as a for-profit venture. The model indicates that this would not produce the correlation between assistance and proliferation behaviors; skill transfers in the absence of a signal incentivize D to make an acceptable offer more often in the interior solution. See Lemma 7 for proof. More proliferation can occur in cases where R’s natural proliferation value is negative, but this does not fit the empirical observations that apply to states broadly.

<sup>7</sup>Similar results would follow if assistance were free or benefited the donor by providing improving knowledge in the other direction. The only difference is that D would make the transfer more often.

<sup>8</sup>Formally, the first derivative of  $\delta(\bullet)$  is strictly positive and its second derivative is strictly negative. Diminishing marginal returns are a natural assumption to make about both speed and costs of proliferation. If those were linear or had increasing marginal returns to skill, as speed approached 1,

and the competent skill type's skill as  $s' > s$ . Then I write the discount rate of a competent type without assistance as  $\delta(s')$  and the discount rate of a competent type with assistance as  $\delta(s' + t)$ . The incompetent type's rates are analogous.

In addition, let  $c(\bullet) \rightarrow (0, \infty)$  map an underlying skill and transfer (or lack thereof) to a cost of proliferation. I assume that the function is strictly decreasing and that skill and aid have diminishing marginal cost reductions.<sup>9</sup> Like before, the cost of a competent type without assistance is  $c(s')$ , the cost of a competent type with assistance is  $c(s' + t)$ , and so forth.

Let  $b > 0$  represent the coercive value, prestige gain, or domestic benefit from obtaining nuclear weapons. Then R's payoff equals  $\delta(\bullet)b - c(\bullet)$ , where each of those functions is defined by the type and assistance choice as described above. Meanwhile, let  $e > 0$  represent the externality D suffers when R proliferates. This incorporates any security losses between rivals, deviations from alliance patron's preferred policy in negotiations between friends, and environmental damage or risk of catastrophic accident. Then D's payoff equals  $\delta(\bullet)e$ .

To recap, the game proceeds as follows:

1. Nature draws R's proficiency level, which is private information to R
2. D chooses whether to provide assistance
3. If D provided assistance, Nature sends it a noisy signal, indicating that R either has "low" or "high" competence
4. D offers concessions to R
5. R accepts those concessions or proliferates, with proliferation cheaper and faster if D provided assistance earlier

Before analyzing the game's solution, a few notes about the game's assumptions are in order. First, under this setup, R automatically receives the transfer if D offers

---

increasing skill further would eventually yield a proliferation time greater than 1—meaning that proliferation is *faster* than instantaneous. Similarly, as costs approached 0, increasing skill further would eventually yield a cost less than 0—meaning that the production of nuclear weapons *makes* money. Decreasing marginal returns allow the speed and cost to approach 1 and 0, respectively, without causing these absurdities.

<sup>9</sup>Formally, the first derivative of  $c(\bullet)$  is strictly negative and its second derivative is strictly positive.

it. One may wonder whether the results would change if it were subject to the same accept/reject decision as the offer. In particular, one may worry that the incompetent type would have incentive to reject to avoid revealing its lack of skill. However, including such a decision does not alter the results. Transfers exhibit an “unraveling” principle. The competent type has nothing to lose in this game by revealing its private information and has skill to gain by receiving the transfer. It therefore would accept. Thus, if the incompetent type declines, it reveals its type by virtue of *not* accepting assistance. This unraveling effect also applies to more complicated setups with more than just a binary type space.

Second, R receives no direct value from the transfer; any benefit comes from extracting a better deal from D later or a more attractive proliferation option. In practice, assistance on nuclear power plants is a tangible benefit, and one that a recipient would not want to see revoked. Nevertheless, I make this assumption to stack the deck against transfers. If donor states could offer less to recipients post-transfer to induce nuclear reticence as a consequence of the existing cooperation, then assistance looks more attractive. It is therefore more surprising to observe transfers in the game presented here.

Third, the model excludes a preventive war decision, which is the focal point of many existing models of nuclear proliferation (Debs and Monteiro, 2014; Spaniel, 2015; Bas and Coe, 2016). Preventive war is a critical feature of some proliferation interactions, but it is not universal. When the United States bargains with allies, for example, policymakers in Washington would rather permit these countries to develop a nuclear weapon than engage in a war that would ruin the friendship. Sworn enemies may also decline preventive action when the target has a large conventional deterrent or the protection of a strong ally (Debs and Monteiro, 2016). Excluding preventive war from the analysis also ensures that the mechanism is not the result of some complicated second-order interaction with it.

## 4 Optimal Bilateral Aid

Because this is an extensive form game of incomplete information, perfect Bayesian equilibrium is the appropriate solution concept. PBE requires that players’ actions are optimal given their beliefs at every stage. It also means that D updates its belief and

conditions its offer based on the signal should it provide assistance.

Throughout, I assume that each type of proliferator has a non-negative value for building nuclear weapons in the absence of a transfer.<sup>10</sup> This is mostly for technical reasons, as it avoids having to deal with corner solutions. The corner solutions behave similarly to the cases I analyze here with a minor exception: D finds transfers more valuable when one type has a negative proliferation payoff but the other one has a positive proliferation payoff. This is because D could still induce the incompetent type to not proliferate without offering any concessions, making the functional cost of acquiring the signal lower. As such, focusing on the interior solution also stacks the deck against a transfer occurring.<sup>11</sup>

In addition, I also focus the analysis on cases where D would prefer brokering a deal with all types regardless of the transfer. That is, the amount of concessions D would have to offer to induce each type to accept is less than D's disutility for suffering proliferation.<sup>12</sup> This automatically holds if states are security rivals and whatever benefits the proliferator receives after acquiring weapons implies an equivalent security loss to its opponent. It also applies to allies when the nonproliferator finds the non-security externalities to be sufficiently high, perhaps due to damage to the regime or the risks of accident. The United States takes observable actions that suggests it fits these cases. For example, it is difficult to explain American concessions to Japan—a close ally and technically proficient country—if policymakers in Washington did not care about these externalities.<sup>13</sup>

Now to solve the game. The interaction hinges on whether D performs better by providing assistance. Although the decision itself is simple, the thought process behind the optimal choice is not. To begin unraveling the logic, if D does not provide assistance, it faces a classic risk-return tradeoff. It must choose from one of two options. First, it can offer a generous settlement designed to induce both competent and incompetent

---

<sup>10</sup>Formally, this requires the incompetent type's proliferation value without a transfer to be at least 0, or  $\delta(s)b - c(s) \geq 0$ .

<sup>11</sup>Another corner solution consists of both types having negative payoffs for proliferation. In this case, the game's solution is trivial: D provides no assistance and offers nothing, and R does not build regardless of its type.

<sup>12</sup>Formally, this requires D's minimum payment to the competent type following a transfer to be less than the time-adjusted externality it suffers, or  $\delta(s' + t)b - c(s' + t) < \delta(s' + t)e$ .

<sup>13</sup>The United States returned Okinawa and offered security guarantees after Japan cemented its non-nuclear status. See Campbell and Sunohara 2004, 222-225.

types to accept. This guarantees that D avoids suffering the externality, but it also means that D overpays the incompetent type for its compliance. Alternatively, D could offer a stingy amount designed to induce only the incompetent type to accept. This implies giving fewer concessions but also results in D suffering the externality whenever R is incompetent.

With that in mind, D's choice depends on its prior belief. Let  $p^* \equiv \frac{\delta(s')e - \delta(s')b + c(s')}{\delta(s')e - \delta(s)b + c(s)}$ . Then the following summarizes D's decision following no assistance:

**Lemma 1.** *(Skeptical) If D believes R is sufficiently likely to be incompetent, it settles with only that type and induces the competent type to proliferate. That is, if  $p > p^*$ , D offers  $\delta(s)b - c(s)$ . The incompetent type accepts and the competent type rejects.*

**Lemma 2.** *(Worried) If D believes R is sufficiently likely to be competent, it settles with both types, and no proliferation occurs. That is, if  $p < p^*$ , D offers  $\delta(s')b - c(s')$ . Both types accept.*

Thus, D's prior belief determines whether proliferation may occur or not. If it is sufficiently skeptical of R's competence, it makes a smaller offer; if it is sufficiently worried about R's competence, it makes a larger offer. Whether D is skeptical or worried has downstream consequences on whether D wants to provide assistance and the empirical implications that come with the transfer.

Nevertheless, some logic applies regardless of whether D is skeptical or worried. Should D provide assistance, the following subgames have similar risk-return tradeoffs to Lemmas 1 and 2. Two important facts make them distinct, however. First, endowing R with greater capacity means that D must make greater concessions to buy off either type. For example, the competent type's payoff for proliferating moves from  $\delta(s')b - c(s')$  before assistance to  $\delta(s' + t) - c(s' + t)$  after. Because the development speeds increase and costs decrease in competency, the latter payoff is larger. Therefore, D must increase its offer to induce the competent type to accept. The same is true for the incompetent type.

More subtly, the changes in proliferation payoffs alter incentives in the risk-return tradeoff. Because competence has diminishing marginal returns, the difference between the competent type's proliferation payoff and the incompetent type's proliferation payoff decreases. This decreases the overpayment to the incompetent type necessary to induce the competent type to accept. In turn, all else equal, the safe offer looks more attractive.

Thus, D must be more convinced that R is incompetent to make the risky offer than before. The appendix shows that critical threshold analogous to Lemmas 1 and 2 is  $r^* = \frac{\delta(s'+t)e - \delta(s'+t)b + c(s'+t)}{\delta(s'+t)e - \delta(s'+t)b + c(s'+t)}$ . This is strictly greater than  $p^*$ .<sup>14</sup>

Second, assistance forces D to update its beliefs. Intuitively, receiving a high signal increases D's belief that it is facing a competent type. All else equal, this makes D more inclined to make the safe offer. Receiving a low signal increases D's belief that it is facing an incompetent type.<sup>15</sup> All else equal, this makes D more inclined to make the risky offer.

From here, D can calculate its expected payoff for giving assistance and compare it to its payoff for making no transfer. This is computationally intensive, so I leave that work to the appendix. However, a couple of theoretical implications follow from the logic developed above.

**Lemma 3.** (*Persuasiveness*) *A necessary condition for D to provide aid in equilibrium is that the set of types D induces to accept with its offers following the signal is not identical to the set of types D induces to accept if it does not receive the signal.*

Put differently, if D provide assistance, one of the signals must be persuasive enough to change which types D buys off. For example, without assistance, imagine that D would make the risky offer that only the incompetent type accepts. However, upon receiving the high signal, D would instead make the safe offer.<sup>16</sup> Then the signal is persuasive because it swayed D's belief enough that offer no longer targets the same type of R. If D would continue offering an amount that just the incompetent type would accept for either signal, then the signal is not persuasive.

Why are only persuasive signals worth purchasing? If the signal does not change D's holistic bargaining strategy, then it provides no real benefit. To the contrary, it reduces D's payoff in two ways. First, D pays the monetary cost of the transfer. Second, as an indirect effect, it forces D to offer more concessions to buy off whichever set of types it

---

<sup>14</sup>For derivation of  $r^*$ , see the proof of Lemmas 1 and 2. For proof that it is greater than  $p^*$ , see Lemma 7 in the appendix. Manipulating parameters to make types behave more similarly commonly results in a higher probability of settlement in other contexts (Reed, 2003; Spaniel, 2018).

<sup>15</sup>See Lemma 6 in the appendix for proof.

<sup>16</sup>Note that the "risky" and "safe" offers differ depending on whether R received assistance. I use these phrases here to describe whom the offer is attempting to appease rather than the size of the offer.

had originally planned to appease. As such, in equilibrium, D only provides assistance to obtain persuasive signals.

**Lemma 4.** (*Responsiveness*) *A necessary condition for D to provide aid in equilibrium is that its optimal offer depends the signal it receives.*

Put differently, if D provides assistance, the offer it makes following the low signal must be different than the offer it makes following the high signal. It cannot be the case that D makes the safe offer that induces both types to accept regardless of the signal received. It also cannot be the case that D makes the risky offer that induces only the incompetent type to accept regardless of the signal received.

The intuition is straightforward. Imagine that D would make the safe offer no matter the signal. Then it must offer a quantity commensurate with the competent type's skill with assistance. However, suppose instead that D did make a transfer. Then D can still buy off both types. But the amount necessary to do that is a quantity commensurate with the competent type's skill *without* assistance. This is strictly less than the first case and is therefore a profitable deviation.<sup>17</sup>

A similar logic holds when D would make the risky offer no matter the signal received. The offer must be a quantity commensurate with the incompetent type's skill with assistance to execute this strategy, and it pays that amount some portion of the time. The remaining portion of the time, D suffers the externality at a speed consistent with a competent type with assistance. However, suppose instead D did not make a transfer. It can still buy off the incompetent type, but now the payment need only be commensurate with the incompetent type's skill *without* assistance. Meanwhile, the portion of the time R is competent and rejects, D suffers the externality at a speed consistent with a competent type without assistance. Both of these outcomes are preferable to the first case, and thus not providing assistance is a profitable deviation.

To further pin down D's decision, note that D's posterior belief that R is incompetent following a low signal must be higher than its prior. Likewise, that belief following a low signal must be lower than its prior. Combining this fact with Lemmas 3 and 4 means that D would only be willing to provide assistance if it would want to make the risky offer in response to a low signal and the safe offer in response to a high signal.

---

<sup>17</sup>Formally, D's utility for buying off both types with assistance is  $\delta(s' + t)b - c(s' + t)$ , whereas D's utility for buying off both types without assistance is  $\delta(s')b - c(s')$ . D's payoff is the negative of the offer, and the former offer is larger than the latter.

This leads to the main description of D's assistance decision:

**Proposition 1.** *Suppose the signal is persuasive and D is responsive to it, as defined by Lemmas 3 and 4. Then D provides assistance if the signal is sufficiently strong compared to the skill transferred and the cost of aid.*

The core intuition behind Proposition 1 is straightforward. If the signal meets the requirements of Lemmas 3 and 4, then providing assistance may be useful. But just because the signal provides decent information does not mean it is worth pursuing. Aid still sacrifices bargaining power and has a direct cost. If the information gains are not worth those prices, then D does not provide assistance. Otherwise, it does.

That straightforward description belies some of the nuance to D's decision. To begin, one may wonder whether a signal exists such that D makes the risky offer following a low signal and the safe offer following a high signal. In fact, there is.<sup>18</sup> One can see by observing that as the signal becomes perfectly informative, D's belief that R is incompetent goes to 1 following the low signal and goes to 0 following a high signal. As a result, a sufficiently informative signal implies that D's two post-signal beliefs straddle  $r^*$ , the cutpoint determining whether D prefers making the safe or risky offer post-transfer.

If the signal is sufficiently strong, then D simply compares its payoffs making and not making the transfer.<sup>19</sup> It then picks the option corresponding to the higher payoff. Increasing the signal quality only increases D's utility, so the first prerequisite for wanting to acquire the signal is not at odds with its value for acquisition. Nevertheless, D may not provide assistance even when the signal is perfectly informative. For example, the transfer may sacrifice so much bargaining power that learning the information is not worth having to make so much greater concessions. The direct monetary cost of the transfer can also make aid suboptimal.

Overall, Proposition 1 explains why donor states provide technical assistance to would-be proliferators. Recasting the transfer as information acquisition generates an intuitive theory. What remains unanswered is how assistance affects proliferation rates. In contrast to before, this is not obvious and depends on D's initial beliefs. I start with cases when D is skeptical:

---

<sup>18</sup>I prove this in the process of proving Proposition 1.

<sup>19</sup>This step is the same regardless of whether D's prior belief makes it skeptical or worried, though the actual calculations are different.

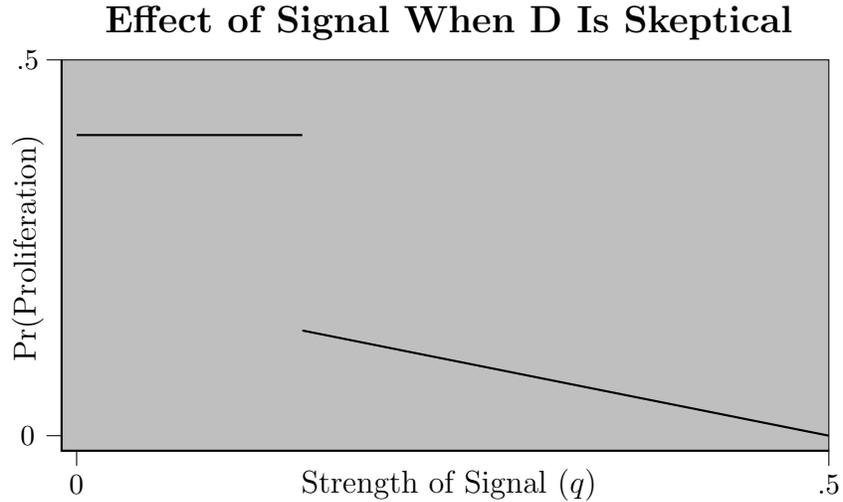


Figure 1: The probability of proliferation when D begins as skeptical. Weak signals are not worth purchasing, so the probability of proliferation remains flat initially. D makes transfers when the signal is sufficiently strong, and the influx of information when D begins receiving the signal causes a discontinuous drop in the probability of proliferation. That probability decreases further as the signal grows stronger.

**Proposition 2.** *Suppose D's prior makes it skeptical of R's competence (i.e.,  $p > p^*$ ). Then the probability of proliferation is strictly lower when D provides assistance. Moreover, the probability of proliferation is weakly decreasing in the strength of the signal.*

Proposition 2 appears to vindicate information theorists. Indeed, if D begins skeptical, observed transfers imply a lower probability of proliferation. Thinking back to Lemma 1 explains why. Under these conditions, in the absence of a transfer, D would make an offer that the competent type would reject. However, if D provides assistance in equilibrium, a high signal must induce D to appease the competent type. Thus, the only time R proliferates is when D wrongly receives the low signal and R is actually competent. This becomes less likely to happen as the signal quality increases.

Figure 1 illustrates the logic. When the strength of the signal is sufficiently weak, D does not provide assistance; the direct cost and the indirect need to give more concessions following a transfer makes the information not worth the price. D always offers an amount tailored to the incompetent type, causing the competent type to reject.

Because D learns no new information, the probability of proliferation remains constant in the strength of the signal, equal to D's prior belief that R is competent.

Nevertheless, further increases to the signal quality make the transfer worthwhile.<sup>20</sup> At that point, the probability of proliferation drops discontinuously because D finally releases all the pent-up signal value.<sup>21</sup> As the signal becomes stronger, the probability it wrongly informs D that R is incompetent decreases. As such, the probability of proliferation decreases continuously from that point forward. When  $q$  approaches .5, the signal becomes perfectly informative. In turn, the probability of miscalculation is 0, and D assuredly reaches a deal with R.

Thus, it appears that assistance has a chilling effect on proliferation. But that is a hasty conclusion:

**Proposition 3.** *Suppose D's prior makes it worried of R's competence (i.e.,  $p < p^*$ ). Then the probability of proliferation is strictly greater when D provides assistance for all noisy signals. The relationship between the strength of the signal and the probability of proliferation is nonmonotonic, with the probability maximizing for middling strengths of the signal for some parameters.*

Proposition 3 demonstrates that additional information does not always facilitate agreement. When D begins as skeptical, its optimal offer guarantees a nonproliferation outcome. As such, any manipulation to the bargaining environment can only increase the probability of proliferation. Here, this happens because the signal may be misleading. A low signal suggests incompetence, nudging D to make an aggressive offer that the competent type would be reject. And while D correctly updates its belief that the incompetent type is more likely than it thought previously, the signal is not perfect. Whatever portion of the time the signal inaccurately indicates low competence, D sees its offer rejected, and proliferation occurs.

Figure 2 illustrates the logic. Like the previous case, D does not provide assistance when the strength of signal is low. As such, the probability of proliferation remains flat in the left side of the figure. In contrast to before, that probability equals 0 because D wants to offer a safe amount given its prior belief.

---

<sup>20</sup>Because increasing signal quality increases the utility for assistance, Proposition 1's condition is not at odds with Lemmas 3 and 4, which also require a sufficiently strong signal to hold.

<sup>21</sup>If no signal is worthwhile, then the probability of proliferation remains flat across the entire range.

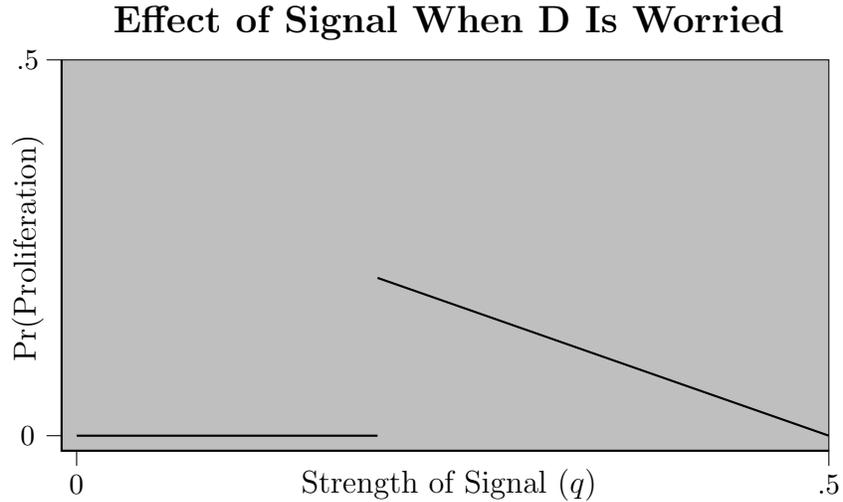


Figure 2: The probability of proliferation when D begins as worried. Weak signals are still not worth purchasing, so the probability of proliferation remains flat initially. When the signal is sufficiently informative that D makes the transfer, the influx of information causes an *increase* in proliferation because it can encourage D to make risky offers. Stronger signals decrease the probability of making a mistake, though.

However, increasing the signal strength beyond a certain point causes D to transfer assistance.<sup>22</sup> This causes a discontinuous *increase* in the probability of proliferation. Specifically, the right side of the graph depicts the probability R is actually a competent type but D falsely receives the low competency signal. That misleading signal causes D to tailor its offer to the wrong type, leading the competent type to reject. The probability of proliferation decreases as the signal becomes stronger, as wrong signals become increasingly less likely. Even so, for this parameter space, the probability of proliferation is strictly greater when D provides assistance than when it does not, unless the signal is *perfectly* informative.

Although Proposition 3 produces a similar empirical implication as the “spreading temptation” logic, my model’s mechanism is distinct. Fuhrmann (2009a) argues that assistance leads to additional proliferation because recipients have lower barriers and quicker paths to nuclear weapons. Within a bargaining game, however, the optimal

<sup>22</sup>As before, D may not want to purchase the signal even if it were perfectly informative. In that case, the probability of proliferation would stay flat at 0 throughout Figure 2.

offer washes out lower costs and faster development speeds.<sup>23</sup> In fact, if skill has diminishing marginal returns, then transfers make D more inclined to pursue safer offers *ceteris paribus*.<sup>24</sup> Instead, the signal can cause D to become more optimistic about R's incompetence, and that belief leads D to pursue riskier offer strategies. Despite the similar empirical implications, exploring my mechanism is critical because the policy recommendations that follow differ. I return to this point in the discussion of transfers and international institutions.

Taking stock, the results depicted in Figures 1 and 2 suggest an empirical conundrum. For one parameter space, countries receiving assistance are less likely to proliferate than those that do not. For the other parameter space, countries receiving assistance are more likely to proliferate than those that do not. One may wonder which scholars would observe more often in practice. Fuhrmann (2009a) indicates the latter, and the model gives theoretical reason to expect this. Consider the cutpoint that determines which parameter space the game falls in:  $p^* = \frac{\delta(s')e - \delta(s')b + c(s')}{\delta(s')e - \delta(s)b + c(s)}$ . Assistance yields more proliferation when  $p$  falls below that cutpoint. Thus, if plausible values of  $p^*$  are high, assistance would correlate with proliferation behaviors.

Although  $p^*$  is a function of many parameters, two salient ones suggest that data would reflect what Fuhrmann (2009a) finds. First, as  $e$  increases,  $p^*$  tends toward 1. Substantively, if the externality of proliferation is large, D must be very sure R is incompetent to make a risky offer. This would be the case for states with deep concerns regarding nuclear accidents and the possibility of catastrophic inadvertent war (Sagan, 1995). Under these circumstances, D takes a safe approach in the absence of a signal. Acquiring the signal can lead to a false sense of skepticism in R's ability, which in turn yields a risky offer and some probability of rejection.

Second,  $p^*$  also increases as  $b$  decreases. Substantively, if the benefit of proliferation is small, the size of offers competent and incompetent types become increasingly similar. In turn, D must be very sure R is incompetent to make the risky offer—which only gives a small bit less than the safe offer—that might cause it to suffer the externality. This

---

<sup>23</sup>A complete information game clarifies this. If D knows that R's competence is  $s$ , then its optimal offer is  $x = \delta(s) - c(s)$ . If R's competence shifts to  $s + t$  because of a transfer, then the optimal offer increases to  $x = \delta(s + t) - c(s + t)$ . In either case, R accepts, and the probability of proliferation remains static at 0. The only way the transfer could result in proliferation is if D prefers allowing proliferation to occur over making the necessary additional concessions. Of course, this point is moot in a complete information game, as D has no incentive to provide assistance if it has nothing to learn.

<sup>24</sup>See Lemma 7 in the appendix.

would be the case if nuclear weapons provide little coercive value and instead only give their possessors deterrent power (Sechser and Fuhrmann, 2017). Like before, D would naturally want to take the safe approach in the absence of a signal. The signal can then give the wrong message and cause D to switch to the risky offer, which may fail.

## 5 Optimal International Assistance

As previewed earlier, institutions also give nuclear assistance. A natural question to ask is whether institutions behave the same way as the donor state in the previous model. I address that now.

Incentive-wise, institutions differ from states in a key way: they do not make coercive policy concessions to would-be proliferators. To illustrate, the United States kept troop deployments and military aid high to South Korea to convince Park Chung-hee to not develop nuclear weapons. In more recent years, the United States relaxed economic policy and began a diplomatic thaw with Iran to convince leaders in Tehran to freeze their nuclear program. The IAEA does not have the ability or means to do this, nor are potential proliferators desperate to convince the IAEA to make bold policy concessions. In short, the IAEA shares the desire to limit proliferation, and it can also make nuclear transfers and read signals about a state's skill. But it does not provide the concessions necessary to convince a state to forgo nuclear weapons.

How does being divorced from coercive bargaining affect an institution's decision to provide assistance? Retooling the model provides some answers. The bargaining subgames remain identical to before. However, rather than D choosing whether to provide assistance, an **O**rganization plays that role. Like D, O does not know R's competency at that time. If it makes the transfer, D still receives a noisy signal. One could conceptualize this as the IAEA receiving verifiable information that they could pass along to D. Play proceeds from there as before.

O's payoff is simple. It wants to minimize proliferation, and providing assistance is costly. To capture that, let 1 be O's standardized value for a nonproliferation outcome and 0 be its standardized value for a proliferation outcome. In addition, O pays  $k$  if it provides assistance.

The bargaining logic D faces remains the same as before. It simply looks at its belief—whether that be the prior or the posterior if it received a signal—and checks

whether making the risky or safe offer is preferable. Meanwhile, O earns its best possible payoff when the game ends without proliferation and it does not pay for the transfer. This leads to a straightforward proposition:

**Proposition 4.** *Suppose D's prior makes it worried of R's competence (i.e.,  $p < p^*$ ). Then O does not provide a transfer regardless of cost. D offers a safe amount, and both types of R accept.*

Intuitively, if O maintains the default bargaining environment, D's optimal offer guarantees a nonproliferation outcome. Any assistance O might provide could only ruin that outcome and come at a direct cost. Able to obtain its favorite outcome by doing nothing, O does just that.

Despite Proposition 4's straightforward logic, it already reveals a key difference between O's and D's incentives. When D is worried about R's competence, it may provide assistance to test the waters. If the signal indicates competence, D makes the safe offer as it would have planned to do otherwise. But if the signal indicates incompetence, D makes the risky offer that results in proliferation whenever the signal provided inaccurate information. Thus, assistance leads to more proliferation in this case with D as the donor. The same is true if O would provide assistance. However, realizing that, O simply chooses not to purchase information that could only sabotage the bargaining outcome.

The institution faces a richer strategic problem when D begins skeptical of R's competence. In the absence of a signal here, D would make an offer that results in proliferation with positive probability. Lemma 3's persuasiveness requirement still holds. That is, for O to want to provide assistance, the signal must change the set of types that D reaches an agreement with. One way this can happen is if the low signal causes D to maintain its risky offer strategy but a high signal induces D to switch to the safe offer strategy:

**Proposition 5.** *Suppose D's prior makes it sufficiently skeptical of R's competence (i.e.,  $p > r^*$ ). Then the probability of proliferation is strictly lower when O provides assistance. Moreover, the probability of proliferation is weakly decreasing in the strength of the signal.*

The logic matches that of Proposition 2. Receiving the high signal causes D to settle with the competent type. Thus, the only way proliferation occurs is if the signal

inaccurately reports incompetency. This probability decreases as the signal becomes stronger. As a result, O's value for assistance is more likely to exceed the costs when the signal is more informative.

Overall, the probability of proliferation matches Figure 1 with a critical exception. A barrier to D providing assistance was that doing so forced D to make additional concessions during the bargaining phase. Here, however, O does not have to pay for those concessions. Thus, it faces one less disincentive for providing assistance. In turn, the minimum strength of signal necessary to make the transfer is strictly lower for O than it is for D.

A close examination of Propositions 4 and 5 reveal that they do not cover middle cases when D's prior falls between  $p^*$  and  $r^*$ . The first condition means that D would want to make a risky offer given its prior and R's baseline possible levels of skill. The second condition means that D would want to make a safe offer given R's possible level of skills post-transfer *if it maintained its prior*.

As explained before, these are not identical cutpoints because assistance makes the types behave more similarly, which disincentivizes D from making risky offers. Lemma 4's responsiveness requirement meant that D would never purchase a signal that triggers an identical offer regardless of what it says. For example, if making a risky offer without a signal was preferable to making the safe offer, then making the safe offer under worse circumstances would yield an even lower utility. But O does not care about forcing D to provide greater concessions. This means that O is willing to make a transfer that switches D to wanting to always make the risky offer to wanting to always make the safe offer.

However, whether D wants to make the safe offer regardless of its signal depends on the strength of the signal. A more powerful signal actually *hurts* O here because a strong low indicator induces D to make the risky offer; if the signal is wrong, then proliferation occurs. This has two drawbacks for O. Directly, the transfer becomes less valuable because it no longer solves the problem with certainty. Indirectly, because the transfer is less valuable, O may not want to provide it at all anymore. This puts the probability of proliferation back at the default level when D is skeptical. The only good news here is that, like Figure 2, both these problems go away as the signal becomes perfectly informative.

The following proposition summarizes this logic:

**Proposition 6.** *Suppose  $D$ 's prior about  $R$ 's competency falls in a middle range (i.e.,  $p \in (p^*, r^*)$ ). Then the probability of proliferation is strictly lower when  $O$  provides assistance. However, the probability of proliferation is nonmonotonic in the strength of the signal.*

It is worth noting that the result on sufficiently strong signals in this parameter space is sensitive to the assumption that the signal is publicly observable. Alternatively, one could conceive of signals that only  $O$  observes that it could then share with  $D$ . In that case,  $O$  would want to conceal the signal regardless of whether it reports low or high competency.  $D$  would not update its belief under these circumstances, and its prior belief induces it to make the safe offer.  $O$  could then obtain a zero probability of proliferation.

Taking stock, the apparent connection between institutional assistance and proliferation behaviors appears more perverse than with bilateral assistance. Here, situations where no proliferation would occur never receive assistance. This gives a false impression that the transfers are unhelpful. The middle region covered in Proposition 6 would counterbalance this, where assistance leads to a zero probability of proliferation. However, IAEA technical cooperation projects are generally small. The difference between  $p^*$  and  $r^*$  may therefore be narrow under these circumstances, meaning that they are not empirically common.<sup>25</sup> The remaining cases feature both assistance and positive probabilities of proliferation, even though the net effect of assistance reduces those chances.

## 6 Conclusion

Why do proliferation opponents give technical assistance to would-be nuclear states? This paper tackled that question from an informational perspective. Convincing states to forgo nuclear weapons requires giving concessions commensurate with their value for proliferation. Opponents cannot easily observe a state's underlying nuclear proficiency, which determines that value. They therefore provide nuclear assistance to improve their estimates. Although such aid forces opponents to make deeper concessions later, the price is sometimes worth the benefit.

---

<sup>25</sup>The average technical cooperation project only has a budget of \$60,000 (Brown and Kaplow, 2014, 424).

The theory helps explain a curious empirical finding in the proliferation literature. States receiving nuclear cooperation agreements from their peers or technical cooperation from the IAEA are more likely to engage in proliferation behaviors. A first-cut interpretation of these results might suggest that assistance backfires. However, the model shows the relationship may not be straightforward. In some cases, assistance helps; in other cases, it increases proliferation rates while simultaneously increasing the proliferation opponent's welfare. Even an international institution that only wants to minimize proliferation may appear to induce it, as it has little incentive to provide assistance to states that will likely reach a deal in the absence of help.

Two policy implications follow from these results. First, although this analysis assumed that donors and international institutions were strategic in their assistance decisions, some may not be in practice. The IAEA, for example, has little oversight over technical cooperation proposals (Brown and Kaplow, 2014, 406), making them susceptible to manipulation by enterprising potential proliferators. Nevertheless, strategic institutions can reduce proliferation by targeting assistance at states with skeptical bargaining partners; blanket information provision can backfire. Bargaining partners can also improve their welfare with careful assistance decisions, though such cases can cause more proliferation.

Second, opponents of potential proliferators should be careful in their donations to international institutions. Opponents and states do not always have aligned preferences in nonproliferation negotiations. Additional skill forces opponents to pay more to induce nuclear reticence. Institutions do not have to increase their concessions to match competence, but opposing states do. Thus, if states wish to work through institutions in this regard, they ought to consider whether the institution would want to fully follow through on future transfers plans.

## 7 Appendix

I begin with some preliminaries, first by pinning down D's beliefs after it gives aid and receives a signal:

**Lemma 5.** *Upon receiving the low signal, D's posterior belief that R is incompetent is  $\frac{p(.5+q)}{.5-q+2pq}$ . Upon receiving the high signal, D's posterior belief that R is incompetent is  $\frac{p(.5-q)}{.5+q-2pq}$ .*

*Proof.* This is a simple application of Bayes' rule. Recall that the prior belief that R is incompetent is  $p$ . D can receive a low signal in two ways: R is actually incompetent and it received the correct signal with probability  $.5 + q$ , or R is actually competent and it received the wrong signal with probability  $.5 - q$ . Therefore, the probability of R being incompetent upon receiving the incompetent signal is:

$$\frac{p(.5 + q)}{p(.5 + q) + (1 - p)(.5 - q)} = \frac{p(.5 + q)}{.5 - q + 2pq} \equiv \underline{r}$$

Likewise, D can receive a high signal in two ways: R is actually incompetent and it received the incorrect signal with probability  $.5 - q$ , or R is actually competent and it received the correct signal with probability  $.5 + q$ . Therefore, the probability of R being incompetent upon receiving the competent signal is:

$$\frac{p(.5 - q)}{p(.5 - q) + (1 - p)(.5 + q)} = \frac{p(.5 - q)}{.5 + q - 2pq} \equiv \bar{r}$$

□

**Lemma 6.** *D's posterior belief that R is incompetent increases following an incompetent signal and decreases following a competent signal (i.e.,  $\underline{r} > p$  and  $\bar{r} < p$ ).*

*Proof.* This requires simple examination of  $\underline{r}$  and  $\bar{r}$ . Comparing  $\underline{r}$  to  $p$  yields:

$$\frac{p(.5 + q)}{p(.5 + q) + (1 - p)(.5 - q)} > p$$

$$p < 1$$

This is true.

Meanwhile, comparing  $\bar{r}$  to  $p$  yields:

$$\frac{p(.5 - q)}{.5 + q - 2pq} < p$$

$$p > 0$$

This is also true.

□

I am now ready to prove the main results.

## 7.1 Proof of Lemmas 1 and 2

Suppose D did not provide assistance. Then the incompetent type accepts if  $x \geq \delta(s)b - c(s)$ , and the competent type accepts if  $x \geq \delta(s')b - c(s')$ .<sup>26</sup> D's optimal offer must therefore be either  $\delta(s)b - c(s)$  or  $\delta(s')b - c(s')$ . Any more than  $\delta(s')b - c(s')$  is a needless concession, any less than  $\delta(s)b - c(s)$  is worse than buying off at least one type, and anything in between is worse than something slightly smaller, which still yields acceptance from the incompetent type and rejection from the competent type.

We can find D's optimal choice by comparing its utilities for each of these choices. Offering  $\delta(s')b - c(s')$  induces acceptance from both types. D therefore receives the negative of that amount as its payoff. Offering  $\delta(s)b - c(s)$  induces the incompetent type to accept but the competent type to reject. In the former case, D receives the negative value of that amount. In the latter case, D suffers its externality at the competent type's pace. Taking the expectation and comparing utilities, the risky offer is better if:

$$-p(\delta(s)b - c(s)) - (1 - p)\delta(s')e > -(\delta(s')b - c(s'))$$

$$p > \frac{\delta(s')e - \delta(s')b + c(s')}{\delta(s')e - \delta(s)b + c(s)} \equiv p^*$$

By analogous argument, the safe offer is better if  $p < p^*$ . □

## 7.2 Optimal Strategies of the Post-Assistance Subgame

The next task is to prove Lemmas 3 and 4. However, these first require analysis of the post-assistance subgame and some basic comparisons between the two options. The assistance subgame is straightforward to prove given Lemmas 1 and 2. If D makes the transfer, it updates its belief according to Bayes' rule as given by Lemma 5. The cost of the transfer is sunk and has no bearing on D's future decisions. As a result, D's optimization problem is the same as without the transfer, except that the two possible competence levels are  $s + t$  and  $s' + t$ . Thus, D prefers the risky post-transfer offer of  $\delta(s + t)b - c(s + t)$  after receiving the low signal if:

$$r > \frac{\delta(s' + t)e - \delta(s' + t)b + c(s' + t)}{\delta(s' + t)e - \delta(s + t)b + c(s + t)} \equiv r^*$$

---

<sup>26</sup>For the usual reasons, no equilibria exist in which a type accepts with any probability less than 1 when indifferent.

By analogous argument, D makes the safe post-transfer offer of  $\delta(s' + t)b - c(s' + t)$  after receiving the low signal if  $\underline{r} < r^*$ . And likewise, after receiving the high signal, D makes the aforementioned risky offer if  $\bar{r} > r^*$  and makes the aforementioned safe offer if  $\bar{r} < r^*$ .

Before moving forward, the following lemma will prove useful:

**Lemma 7.** *After providing aid, the minimum belief that R is the incompetent type for D to make the riskier offer is higher than the analogous minimum belief having not provided aid. Equivalently,  $p^* < r^*$ .*

*Proof.* It is easier to demonstrate this with some restructuring of the problem. First, redefine R's payoff for proliferation as  $v(\bullet) = \delta(\bullet)b - c(\bullet)$ . Rewriting  $r^*$  in this manner generates  $\frac{\delta(s'+t)e - v(s'+t)}{\delta(s'+t)e - v(s+t)}$ . I can prove Lemma 7 by showing that this value is strictly increasing in  $t$ . Taking the derivative with respect to  $t$  yields:

$$\frac{[\delta(s' + t)e - v(s + t)](\delta'(s' + t)e - v'(s' + t)) - [\delta(s' + t)e - v(s' + t)](\delta'(s' + t)e - v'(s + t))}{(\delta(s' + t)e - v(s + t))^2}$$

This is strictly positive if:

$$(\delta(s' + t)e - v(s + t))(\delta'(s' + t)e - v'(s' + t)) > (\delta(s' + t)e - v(s' + t))(\delta'(s' + t)e - v'(s + t))$$

The first half of the left side is larger than the first half of the right side. Therefore, the inequality holds if the second half of the left side is greater than the second half of the right side. Writing this out:

$$\delta'(s' + t)e - v'(s' + t) > \delta'(s' + t)e - v'(s + t)$$

$$v'(s + t) > v'(s' + t)$$

Each component of R's utility function is concave, so the  $v(\bullet)$  is concave. Therefore, this holds.  $\square$

### 7.3 Proof of Lemma 3

In equilibrium there are two possible sets of types that settle following no transfer: (1) all types and (2) just the incompetent type. Consider the first case. Then D's equilibrium utility for making the associated offer is  $-p(\delta(s)b - c(s)) - (1-p)\delta(s')e$ . If D settles with identical types following the transfer, the expected probabilities remain the same. The equilibrium concession to the incompetent type goes up, however, and the competent type develops nuclear weapons faster. Thus, D's equilibrium utility equals  $-p(\delta(s+t)b - c(s+t)) - (1-p)\delta(s'+t)e$ . Each constituent component is smaller in the second case, and D cannot provide assistance that induces those bargaining strategies post-transfer.

In the second case, D's equilibrium utility for making the associated offer is  $-(\delta(s') - c(s'))$ . If it makes the transfer and settles with both types, its utility equals  $-(\delta(s'+t) - c(s'+t))$ . Each constituent component is again smaller in the second case, so D cannot provide assistance that induces those bargaining strategies post-transfer either.  $\square$

### 7.4 Proof of Lemma 4

If D's post-transfer strategy is not a function of its signal, its offer is either consistently (1)  $\delta(s+t)b - c(s+t)$ , which induces the incompetent type to accept and the competent type to reject, or (2)  $\delta(s'+t)b - c(s'+t)$ , which induces both types to accept. Consider the first case. Then the expected utility for providing assistance is  $-p(\delta(s+t)b - c(s+t)) - (1-p)\delta(s'+t)e$ . But the proof for Proposition 3 showed that the utility for not providing assistance and offering  $\delta(s)b - c(s)$  is strictly greater, so D would have profitable deviation.

In the second case, D's expected utility for providing assistance is  $\delta(s'+t)b - c(s'+t)$ . But the proof for Proposition 3 also showed that the utility for not providing assistance and offering  $\delta(s')b - c(s')$  is strictly greater, so again D would have a profitable deviation.  $\square$

### 7.5 Proof of Proposition 1

Lemmas 3, 4, and 6 imply that any signal D opts to acquire in equilibrium must set  $\bar{r} < r^*$  and  $\underline{r} > r^*$ . Substantively, this means that D makes the safe offer following the

high competence signal and the risky offer following the low competence signal.

I can now prove Proposition 1 through three cases. Two points are sufficient for D to provide assistance for all parameter spaces: (1) a high signal must induce the safe offer and the low signal must induce the risky offer and (2) the utility for providing the assistance given the induced strategies exceeds the utility for not.

This leads to three cases. First, suppose  $p < p^*$ . Because  $\bar{r} < p < p^* < r^*$ , D makes the safe offer following a high signal regardless of the signal's strength. In contrast, if  $q = 0$ ,  $\underline{r} = p < p^* < r^*$ . Thus, this is the binding constraint. For notational simplicity, let  $\frac{n}{d} = r^*$ , with  $d > n > 0$ . The following finds the  $q$  necessary to obtain  $\underline{r} > r^*$ :

$$\frac{p(.5 + q)}{.5 - q + 2pq} > \frac{n}{d}$$

$$q(n + pd - 2pn) > \frac{n - pd}{2}$$

Solving for  $q$  without flipping the inequality requires  $n + pd - 2pn > 0$ . This reworks to  $n + pd > pn + pn$ . The first term of the left hand side is larger than the first term of the right hand side, and the second term of the left hand side is larger than the second term of the right hand side. Thus, the  $q$  necessary to obtain  $\underline{r} > r^*$  is:

$$q > \frac{n - pd}{2(n + pd - 2pn)} \quad (1)$$

Because the denominator is positive, the value is positive if  $n - pd > 0$ . This rewrites to  $p < \frac{n}{d} = r^*$ , which is true for this parameter space. It is also strictly less than .5 because  $\frac{n - pd}{2(n + pd - 2pn)} < \frac{1}{2}$  reduces to  $d > n$ .

So if  $q < \frac{n - pd}{2(n + pd - 2pn)}$ , D does not make the transfer. If  $q > \frac{n - pd}{2(n + pd - 2pn)}$ , then it must check its utilities. If D does not provide assistance, it makes the safe offer, and its utility equals  $-(\delta(s')b - c(s'))$ . If it provides assistance, its payoff is more complicated. With probability  $p(.5 - q) + (1 - p)(.5 + q) = .5 + q - 2pq$ , it receives the high competence signal.<sup>27</sup> Following the high signal, it offers  $\delta(s' + t)b - c(s' + t)$ , and both types accept. With probability  $p(.5 + q) + (1 - p)(.5 - q) = .5 - q + 2pq$ , it receives the low signal.<sup>28</sup> Following the low signal, it offers  $\delta(s + t)b - c(s + t)$ . Only the incompetent type

---

<sup>27</sup>The first term is the probability of receiving that signal and R is incompetent, and the second term is the probability of receiving that signal and R is competent.

<sup>28</sup>The first term is the probability of receiving that signal and R is incompetent, and the second term is the probability of receiving that signal and R is competent.

accepts. Conditional on receiving the low signal, this occurs with probability  $\frac{p(.5+q)}{.5-q+2pq}$ . The competent type rejects, giving D a payoff of  $-\delta(s' + t)e$ . Conditional on receiving the low signal, this occurs with probability  $\frac{(1-p)(.5-q)}{.5-q+2pq}$ .

Putting all this together, D prefers providing assistance if:

$$\begin{aligned}
& -(.5 - q + 2pq) \left( \frac{p(.5 + q)}{.5 - q + 2pq} (\delta(s + t)b - c(s + t)) + \frac{(1 - p)(.5 - q)}{.5 - q + 2pq} (\delta(s' + t)e) \right) \\
& - (.5 + q - 2pq)(\delta(s' + t)b - c(s' + t)) - k > -(\delta(s')b - c(s'))
\end{aligned}$$

Rearranging to help solve for  $q$  yields:

$$\begin{aligned}
& q[(1 - p)(\delta(s' + t)e) - p(\delta(s + t)b - c(s + t)) - (1 - 2p)(\delta(s' + t)b - c(s' + t))] \\
& > .5(\delta(s' + t)b - c(s' + t)) + .5p(\delta(s + t)b - c(s + t)) + .5(1 - p)(\delta(s' + t)e) \\
& \quad - (\delta(s')b - c(s')) + k
\end{aligned}$$

To solve for  $q$  without flipping the inequality requires ensuring that the non- $q$  term on the left hand side is strictly greater than 0. I show this by instead demonstrating that  $(1 - p)(\delta(s' + t)e) - p(\delta(s + t)b - c(s + t)) - (1 - 2p)(\delta(s' + t)b - c(s' + t))$  is greater than 0, as this is less than the quantity in question.

$$(1 - p)(\delta(s' + t)e) - p(\delta(s + t)b - c(s + t)) - (1 - 2p)(\delta(s' + t)b - c(s' + t)) > 0$$

$$\delta(s' + t)e > \delta(s' + t)b - c(s' + t)$$

This is true. Therefore, the utility for providing aid exceeds the utility for not if:

$$q > \frac{(1 - p)(\delta(s' + t)e) - p(\delta(s + t)b - c(s + t)) - (1 - 2p)(\delta(s' + t)b - c(s' + t))}{.5(\delta(s' + t)b - c(s' + t)) + .5p(\delta(s + t)b - c(s + t)) + .5(1 - p)(\delta(s' + t)e) - p(\delta(s)b - c(s)) - (1 - p)(\delta(s')e) + k} \quad (2)$$

In turn, D provides assistance if  $q$  exceeds the values in both Lines 1 and 2. Note that although the  $q$  condition for the necessary beliefs is constrained between 0 and .5, the right hand side of the Line 2 may exceed .5. If this is the case, then D does not

provide aid regardless of  $q$ .

Second, suppose  $p \in (p^*, r^*)$ . In this case, D makes the risky offer in the absence of assistance. Just like before, because  $p < r^*$ , any high signal makes  $\bar{r} < r^*$ . Meanwhile, Line 1 already solves for the minimum necessary signal to obtain  $\underline{r} > r^*$ .

Thus, the only difference is in the utility comparison. The utility for making the transfer is the same as in the first case. But now, the value for not making the transfer is  $-p(\delta(s)b - c(s)) - (1 - p)(\delta(s')e)$ . D therefore prefers making the transfer if:

$$\begin{aligned}
 & -(.5 - q + 2pq) \left( \frac{p(.5 + q)}{.5 - q + 2pq} (\delta(s + t)b - c(s + t)) + \frac{(1 - p)(.5 - q)}{.5 - q + 2pq} (\delta(s' + t)e) \right) \\
 & - (.5 + q - 2pq)(\delta(s' + t)b - c(s' + t)) - k > -p(\delta(s)b - c(s)) - (1 - p)(\delta(s')e)
 \end{aligned}$$

Rearranging yields:

$$\begin{aligned}
 & q[(1 - p)(\delta(s' + t)e) - p(\delta(s + t)b - c(s + t)) - (1 - 2p)(\delta(s' + t)b - c(s' + t))] \\
 & > .5(\delta(s' + t)b - c(s' + t)) + .5p(\delta(s + t)b - c(s + t)) + .5(1 - p)(\delta(s' + t)e) \\
 & \quad - p(\delta(s)b - c(s)) - (1 - p)(\delta(s')e) + k
 \end{aligned}$$

The term next to  $q$  is identical to the first case and is therefore greater than 0; I do not need to flip the inequality to solve for  $q$ . Doing that gives:

$$q > \frac{(1 - p)(\delta(s' + t)e) - p(\delta(s + t)b - c(s + t)) - (1 - 2p)(\delta(s' + t)b - c(s' + t))}{.5(\delta(s' + t)b - c(s' + t)) + .5p(\delta(s + t)b - c(s + t)) + .5(1 - p)(\delta(s' + t)e) - p(\delta(s)b - c(s)) - (1 - p)(\delta(s')e) + k} \quad (3)$$

In turn, D provides assistance if  $q$  exceeds the values in both Lines 1 and 3. Like before, the right hand side of Line 3 may exceed .5. If this is the case, then D does not provide aid regardless of  $q$ .

Third, suppose  $p > r^*$ . Like the second case, D makes the safe offer in the absence of assistance. Thus, Line 3 already solves for the minimum necessary signal to obtain the utility requirement. Unlike either of the previous two cases, because  $\underline{r} > p > r^*$ , any low signal induces D to make the risky offer. Now the signal must be sufficiently strong to force  $\bar{r} < r^*$ . The following finds the  $q$  necessary to obtain  $\bar{r} < r^*$ :

$$\frac{p(.5 - q)}{.5 + q - 2pq} < \frac{n}{d}$$

$$q(n + pd - 2pn) > \frac{pd - n}{2}$$

Solving for  $q$  without flipping the inequality requires  $n + pd - 2pn > 0$ , which the first case already demonstrated was true. Thus, for  $q$  to fulfill the requirement:

$$q > \frac{pd - n}{2(n + pd - 2pn)} \quad (4)$$

Because the denominator is positive, the value is positive if  $pd - n > 0$ . This rewrites to  $p > \frac{n}{d} = r^*$ , which is true for this parameter space. It is also less than .5 because  $\frac{pd - n}{2(n + pd - 2pn)} < \frac{1}{2}$  reduces to  $p < 1$ .

In turn, D provides assistance if  $q$  exceeds the values in both Lines 3 and 4. Once more, the right hand side of Line 3 may exceed .5. If this is the case, then D does not provide aid regardless of  $q$ .  $\square$

## 7.6 Proof of Proposition 2

The proof follows easily from earlier results. If  $p > p^*$  and D does not provide assistance, it makes the risky offer. The incompetent type accepts, and the competent type rejects. The probability of proliferation is  $1 - p$ . With assistance, the probability of proliferation is the probability that R is competent *and* the probability D receives a low signal given that R is competent. This is  $(1 - p)(.5 - q)$ , which is strictly less than  $1 - p$ .

All that is left to show is that the probability of proliferation is monotonic in  $q$ . The proof of Proposition 1 showed that D does not provide assistance if  $q$  is sufficiently low and does if  $q$  is sufficiently high.<sup>29</sup> For  $q$  below the threshold, the probability of proliferation equals  $1 - p$  and is therefore static in  $q$ . After the threshold, the probability is the probability is  $(1 - p)(.5 - q)$ . That probability is strictly decreasing in  $q$ . Moreover, because the cutpoint on  $q$  is strictly greater than 0, the probability of proliferation discontinuously drops at the cutpoint.  $\square$

---

<sup>29</sup>“Sufficiently high” here can exceed .5, in which case the probability of proliferation remains static at  $1 - p$ .

## 7.7 Proof of Proposition 3

This proof also follows easily from earlier results. If  $p < p^*$  and D does not provide assistance, it makes the safe offer. Both types accept, and the probability of proliferation is 0. With assistance, the probability of proliferation is the probability that R is competent *and* the probability D receives a low signal given that R is competent. This is  $(1 - p)(.5 - q)$ , which is strictly greater than 0.

All that is left is to show that the probability of proliferation is nonmonotonic in  $q$ . The proof of Proposition 1 showed that D does not provide assistance if  $q$  is sufficiently low and does if  $q$  is sufficiently high.<sup>30</sup> For  $q$  below the threshold, the probability of proliferation is 0 and therefore static in  $q$ . After the threshold, the probability is  $(1 - p)(.5 - q)$ . That probability is strictly decreasing in  $q$ . However, the probability discontinuously increases at the cutpoint, making the overall relationship nonmonotonic.  $\square$

## 7.8 Proof of Proposition 5

If O does not provide assistance, D makes the risky offer because  $p > r^* > p^*$ . The probability of proliferation is  $1 - p$ , and therefore O's payoff is  $-(1 - p)$ .

If O provides assistance, it must change the probability of proliferation. Otherwise, O's payoff would be  $-(1 - p) - k$ , and it could profitably deviate to no assistance. Because  $p > r^*$ ,  $\underline{r} > r^*$ . In words, receiving the low signal implies that D makes the risky offer. Thus, to change the probability of proliferation,  $\bar{r}$  must be below  $r^*$ . Line 4 showed that this requires  $q > \frac{pd - n}{2(n + pd - 2pm)}$ . If O provides assistance in this case, the probability of proliferation drops to  $(1 - p)(.5 - q)$ . Moreover, this probability is decreasing in  $q$ .

The last things to check are that O would actually want to do this. Unlike D, O's utility for provision is straightforward. Its utility is only a function of the probability of proliferation and its cost of provision, and not a complicated interaction with the concession made. Thus, O provides assistance if:

$$-(1 - p)(.5 - q) - k > -(1 - p)$$

---

<sup>30</sup>Again, "sufficiently high" here can exceed .5, in which case the probability of proliferation remains static at  $1 - p$ .

$$q > \frac{k}{1-p} - .5$$

As such, the probability of proliferation is weakly decreasing in  $q$  because larger values of  $q$  shift the parameters into the region where O provides, the probability of proliferation is discontinuously lower at the cutpoint, and the probability decreases in  $q$  thereafter.  $\square$

## 7.9 Proof of Proposition 6

If O does not provide assistance, D makes the risky offer because  $p > p^*$ , just as in the last case. The probability of proliferation is  $1 - p$ . O could profitably deviate to not providing assistance if the probability of proliferation remained flat after the transfer, so the claim about assistance lowering the probability of proliferation is true.

The more difficult task is to show that the probability of proliferation is nonmonotonic in the strength of the signal. Because  $p < r^*$ ,  $\bar{r}$  must be less than  $r^*$ . Thus, the only complication left is whether  $\underline{r}$  is less than or greater than  $r^*$ . Line 1 gives this cutpoint as  $q = \frac{n-pd}{2(n+pd-2pn)}$ . So if  $q < \frac{n-pd}{2(n+pd-2pn)}$ , the probability of proliferation is 0 following assistance. In turn, O provides assistance if:

$$-k > -(1-p)$$

$$k < 1-p$$

The probability of proliferation is flat in this region.

If  $q > \frac{n-pd}{2(n+pd-2pn)}$ , then the decision problem becomes identical to Proposition 5. The probability of proliferation following a transfer is  $1 - p$ , and O provides assistance if  $q > \frac{k}{1-p} - .5$ . Rearranging in terms of  $k$ ,  $k < (1-p)(q + .5)$ . This is a more stringent requirement than when  $q$  is less than the critical cutpoint. This means that, for some  $k$  values, the probability of proliferation is 0 for low  $q$  values, goes back to  $1 - p$  after it goes above the cutpoint of  $q$ , and then goes back down to  $(1-p)(.5 - q)$  when  $q$  is sufficiently high that  $q > \frac{k}{1-p} - .5$ . The relationship is also monotonic if  $k$  is low enough that O provides assistance regardless, as it goes from 0 to  $(1-p)(.5 - q)$  at the cutpoint and decreases monotonically thereafter.

## References

- Arena, Philip and Scott Wolford. 2012. "Arms, intelligence, and war." *International Studies Quarterly* 56(2):351–365. 3
- Bas, Muhammet A and Andrew J Coe. 2016. "A dynamic theory of nuclear proliferation and preventive war." *International Organization* 70(4):655–685. 5, 10
- Brown, Robert L and Jeffrey M Kaplow. 2014. "Talking peace, making weapons: IAEA technical cooperation and nuclear proliferation." *Journal of Conflict Resolution* 58(3):402–428. 3, 4, 7, 23, 24
- Campbell, Kurt M and Tsuyoshi Sunohara. 2004. Japan: Thinking the Unthinkable. In *The Nuclear Tipping Point: Why States Reconsider Their Nuclear Choices*, ed. Kurt M Campbell, Robert J Einhorn and Mitchell B Reiss. Washington, DC: Brookings Institution Press pp. 218–253. 11
- Debs, Alexandre and Nuno P Monteiro. 2014. "Known unknowns: Power shifts, uncertainty, and war." *International Organization* 68(1):1–31. 10
- Debs, Alexandre and Nuno P Monteiro. 2016. *Nuclear Politics*. New York: Cambridge University Press. 5, 10
- Doyle, James. 2008. Argentina and Brazil. In *Nuclear safeguards, security, and nonproliferation: Achieving security with technology and policy*, ed. James Doyle. Oxford: Butterworth-Heinemann pp. 133–148. 7
- Drezner, Daniel W. 1999. *The sanctions paradox: Economic statecraft and international relations*. Number 65 Cambridge University Press. 6
- Fearon, James D. 1995. "Rationalist explanations for war." *International Organization* 49(3):379–414. 5
- Fuhrmann, Matthew. 2009a. "Spreading temptation: proliferation and peaceful nuclear cooperation agreements." *International Security* 34(1):7–41. 1, 3, 4, 6, 7, 18, 19
- Fuhrmann, Matthew. 2009b. "Taking a walk on the supply side: The determinants of civilian nuclear cooperation." *Journal of Conflict Resolution* 53(2):181–208. 4

- Gonzalez, Abel J. 2001. Medical management of radiation accidents. In *The Safety of Radiation Sources and the Security of Radioactive Materials*, ed. Igor A Gusev, Angela K Guskova and Fred A Mettler. Boca Raton: CRC Press pp. 133–148. 7
- Hersman, Rebecca KC and Robert Peters. 2006. “Nuclear U-turns: Learning from South Korean and Taiwanese rollback.” *Nonproliferation Review* 13(3):539–553. 7
- Hymans, Jacques EC. 2012. *Achieving nuclear ambitions: scientists, politicians, and proliferation*. Cambridge University Press. 7
- Jo, Dong-Joon and Erik Gartzke. 2007. “Determinants of nuclear weapons proliferation.” *Journal of Conflict Resolution* 51(1):167–194. 4
- Keohane, Robert O. 1984. *After hegemony: Cooperation and discord in the world political economy*. Princeton: Princeton University Press. 3
- Kirge, John. 2006. “Atoms for Peace, scientific internationalism, and scientific intelligence.” *Osiris* 21(1):161–181. 7
- Kroenig, Matthew. 2009. “Importing the bomb: Sensitive nuclear assistance and nuclear proliferation.” *Journal of Conflict Resolution* 53(2):161–180. 4
- Kydd, Andrew H. 2010. “Rationalist approaches to conflict prevention and resolution.” *Annual Review of Political Science* 13:101–121. 2
- Miller, Nicholas L. 2014. “The secret success of nonproliferation sanctions.” *International Organization* 68(4):913–944. 5, 6
- Mitchell, Derek J. 2004. Taiwan’s Hsin Chu Program: Deterrence, Abandonment, and Honor. In *The Nuclear Tipping Point: Why States Reconsider Their Nuclear Choices*, ed. Kurt M Campbell, Robert J Einhorn and Mitchell B Reiss. Washington, DC: Brookings Institution Press pp. 293–314. 6
- Montgomery, Alexander H and Adam Mount. 2014. “Misestimation: Explaining US Failures to Predict Nuclear Weapons Programs.” *Intelligence and National Security* 29(3):357–386. 1, 5
- Paul, TV. 2000. *Power versus prudence: Why nations forgo nuclear weapons*. Montreal: McGill-Queen’s University Press. 5

- Reed, William. 2003. "Information, power, and war." *American Political Science Review* 97(4):633–641. 13
- Reiss, Mitchell. 1988. *Without the bomb: The politics of nuclear nonproliferation*. New York: Columbia University Press. 5
- Sagan, Scott Douglas. 1995. *The limits of safety: Organizations, accidents, and nuclear weapons*. Princeton: Princeton University Press. 19
- Schwartz, Stephen I. 2011. *Atomic audit: The costs and consequences of US nuclear weapons since 1940*. Washington, DC: Brookings Institution Press. 5
- Sechser, Todd S and Matthew Fuhrmann. 2017. *Nuclear Weapons and Coercive Diplomacy*. New York: Cambridge University Press. 20
- Singh, Sonali and Christopher R Way. 2004. "The correlates of nuclear proliferation: A quantitative test." *Journal of Conflict Resolution* 48(6):859–885. 4
- Smith, Bradley C and William Spaniel. 2018. "Introducing  $\nu$ -CLEAR: A latent variable approach to measuring nuclear proficiency." *Conflict Management and Peace Science, Forthcoming* . 4
- Spaniel, William. 2015. "Arms negotiation, war exhaustion, and the credibility of preventive war." *International Interactions* 41(5):832–856. 5, 10
- Spaniel, William. 2018. "Only here to help? Bargaining and the perverse incentives of international institutions." *International Studies Quarterly* 62(1):14–22. 13
- Volpe, Tristan A. 2017. "Atomic leverage: Compellence with nuclear latency." *Security Studies* 26(3):517–544. 5