Scientific Intelligence, Nuclear Assistance, and Bargaining*

William Spaniel†

March 11, 2021

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 proficiency. This allows donors to better calibrate other concessions to the recipient. In equilibrium, donors 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 misleading signal indicating incompetence. This paper works toward understanding how scientific intelligence affects international negotiations, an under-explored subject matter for political scientists.

*Forthcoming, *Conflict Management and Peace Science*.

†Department of Political Science, University of Pittsburgh, Posvar Hall 4446, Pittsburgh, PA 15260. (williamspaniel@gmail.com, http://williamspaniel.com). I would like to thank Robert Axelrod, Terry Chapman, James Morrow, Jonathan Markowitz, Jacob Otto, Brad Smith, Alviss Star, Scott Wolford, and the participants of seminars at the University of Michigan, University of Southern California, and University of Texas for their comments on earlier drafts.

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 outward premise was simple. To disincentivize further weapons proliferation, capable states ought to offer peaceful nuclear technology to compliant states. More than a half century later, technology sharing remains a central tenant of the nonproliferation regime.

Nevertheless, nuclear transfers are puzzling. To convince other countries to forgo the bomb, donors make nuclear proliferation cheaper and easier. That powerful countries wish to give concessions to ensure nonproliferation is not surprising. What is surprising is that they use nuclear technology as a tool instead of other economic, 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, donors improve recipients' ability to proliferate and "spread temptation" (Fuhrmann, 2009*a*). This would seem to force donors to give further concessions down the road or result in more countries proliferating outright.

If sharing technology facilitates proliferation, why do donors make such transfers? A couple of existing theories help explain the behavior. One perspective is that integrating nuclear industries helps states monitor the activities of would-be proliferators. The United States, for example, deciphered that Taiwan had begun working toward a bomb by noticing that its requests for nuclear technology exceeded what its civilian research programs could handle (Hersman and Peters, 2006, 544). A second perspective views nuclear assistance as result of Cold War pressures (Medhurst, 1997; Colgan and Miller, 2019). That is, the U.S. was aware of the problem but ignored it in the broader effort to contain the Soviet Union.

I do not dispute that these mechanisms help explain the prevalence of such programs. However, they are not the full story. Historians are clear on this subject: scientific intelligence became a critical need following World War II (Doel, 1997), and transfer programs were a central hub for intelligence gathering for other states' nuclear programs (Krige, 2006). This has been true from the start of such programs and continues to the present, and so I examine how scientific intelligence gathering affects the broader bargaining relationship.

To convince would-be proliferators to end nuclear weapons programs, concerned states need to offer a deal. However, intelligence estimates of foreign proliferation abilities 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. Integrating nuclear programs provides a solution, allowing donors to better understand a recipient's willingness and native proficiency. Learning this information is sometimes worth the drawback of facilitating further proliferation progress.

To develop this argument, I build a model of proliferation, assistance, and uncertainty. A donor does not know a recipient's underlying nuclear preference. It begins by choosing whether to transfer nuclear technology to the recipient. Doing so enhances the recipient's proficiency but also gives the donor an informative—but noisy—signal about the recipient.[1] Afterward, the donor chooses to offer policy concessions. The recipient either accepts or rejects the offer. Rejecting leads to proliferation, which comes cheaper and faster if the recipient received assistance.

The formal analysis indicates that the aforementioned intuition is correct: donors provide assistance to increase their information. However, the model also reveals a counterintuitive implication. 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 begets more proliferation, not less. Why? If assistance were impossible, the donor's initial beliefs might suggest that the recipient is a proficient or willing type, leading to generous offers that are likely to be accepted. In contrast, when making a transfer, the donor may sometimes receive a signal indicating reticence. This can cause the donor to downgrade its offer and lead to proliferation more often if the signal was misleading.

The model 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: a selection problem may drive the data. Assistance can occur when proliferation

[1]Uncertainty over proficiency is a source of *incomplete* information. This makes the model distinct from discussions of imperfect information and informative signals about the proliferator's past actions (Debs and Monteiro, 2014; Bas and Coe, 2016; Miller, 2017; Spaniel, 2019).

is likely in its absence. Despite transfers reducing the barriers to development, a blanket ban on assistance could counterintuitively exacerbate the problem.

From a formal perspective, this paper is closest to Arena and Wolford (2012), which 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. My work differs on two dimensions. First, the substantive issues are distinct. Second, my intelligence gathering mechanism improves the *opponent's* outside option, making it unclear whether the donor would still want to acquire the information.

2 Motivation

Developing a coherent model first requires motivating some of its key assumptions. I do that here, focusing on bargaining, uncertainty, and signals.

2.1 Nuclear Assistance and Asymmetric Information

Nuclear assistance is common across the world, though the scope and donor processes vary from project to project. Nuclear Cooperation Agreements (NCAs) receive the most attention. 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). Cooperation also funnels through the International Atomic Energy Agency. In fact, one of its founding principles is to integrate the nuclear science community and facilitate the exchange of information and training of personnel.[2]

These projects do not intend to assist countries in developing nuclear weapons.[3] However, nuclear technology for civilian applications often has dual uses. For example, assisting a country with its nuclear power seems benign. But work in that area can

[2] See Article III A 3-4 and Article VIII of the IAEA Statute.

[3] This is the realm of "sensitive" nuclear assistance (Kroenig, 2009). It is less surprising that such transfers increase rates of proliferation because the donor is often actively helping in the weaponization effort. In contrast, in my model, the donor is directly harmed if the recipient proliferates and receives no direct utility benefit from giving assistance. These scope conditions distinguish my work from sensitive cases.

train scientists in basic nuclear principles that they could use to construction a weapon. As a result, assistance facilitates building nuclear weapons.

The quantitative "supply side" literature indicates that this is a recipe for more proliferation. Countries with cheaper and faster routes to nuclear weapons seem to be more likely proliferators. The data reflect this. Countries with higher industrial infrastructure pursue and acquire nuclear weapons more frequently (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: why supply assistance given the perverse incentive? A second literature hints at 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 nonproliferation inducements (Reiss, 1988; Paul, 2000; Bas and Coe, 2016; Debs and Monteiro, 2016; Volpe, 2017; Spaniel, 2019; Mehta, 2020). Assistance or not, nuclear weapons are expensive (Schwartz, 2011). Some states nevertheless find nuclear weapons attractive because the benefits outweigh the the costs.

However, proliferation is not inevitable. 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 policy. Both types of states suffer the damage to the nonproliferation regime and the risks of an 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.

Even so, the existence of mutually beneficial agreements does not guarantee their implementation. Particularly problematic for nuclear negotiations is uncertainty. The United States, despite pouring substantial sums of money into intelligence programs, has a poor track record in predicting other states' behaviors (Montgomery and Mount, 2014)—and would conceivably be even worse absent those intelligence interventions. Much of 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 estimate crude timetables. Willingness is another factor. Even if a state has the ability to produce nuclear weapons, it might not if the political will to do so is not present.

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. This is well-established in negotiations between antagonists (Powell, 1999, 83). However, it also applies to negotiations between allies, where patron states do not want to give more weapons than are necessary and thereby embolden protégés to engage in risky behavior (Benson, 2012). Finding this perfect amount is not an easy task. Deals must be commensurate with the speed and cost of nuclear weapons. Thus, if the nonproliferator is uncertain of its opponent's proficiency or willingness, its offer may be insufficient.

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 harm the nonproliferator's bargaining position. After all, regardless of the state's native skill, transfers 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 U.S. fuels them (Mitchell, 2004, 305). Washington 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 lower proliferation barriers.

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, 2009*a*, 20). In the following years, the United States attempted to convince Pakistan to forgo nuclear weapons with large offers of military aid. Unlike South Korea, though, these efforts failed.

2.2 Assistance as Scientific Intelligence

The above concerns lead to the fundamental question this paper addresses: if transfers have downside risks, why do states provide nuclear assistance at all? In fact, assistance has an unappreciated secondary effect: it integrates the donor's nuclear field with the recipient's. To resolve the aforementioned information problem, the U.S. identified human intelligence as a major part of the solution.[4] Although political scientists have not yet examined this, science historians have demonstrated that nuclear assistance exchanges were designed in part to extract information. I now recount some of the historical research on the subject, which will motivate the model in the next section.

As a preface, it would be surprising to find direct evidence that policymakers use nuclear assistance to gather intelligence. This is for two reasons. First, much of the documentation would be locked away under classification (Doel, 1997, 222–224). Examples of direct scientist-to-scientist information disclosure qualify as human intelligence. U.S. law classifies the source and methods of such intelligence acquisition at 50 years and is renewable if the source or method is still in practice, which is likely for nuclear cooperation missions. Second, the United States publicly treats technology transfer as a major soft power initiative, linking scientists around the world with the sophistication of American nuclear physicists. Disclosing an ulterior motive would run contrary to that public relations platform. Indeed, early efforts stalled when a classification failure publicized the CIA's intention to use scientists for espionage (Lexow, 1966, 23).

Nevertheless, a deeper dive into the origins of scientific intelligence paints a clear picture. From the start, the U.S. saw nuclear assistance as a part of a larger intelligence

[4] Much of the substantive evidence here comes from the United States. Part of this is because the records on the U.S. are more open than other countries. But another key aspect is that nuclear proliferation is a "public bad". Hegemonic stability theory predicts that countries that disproportionately benefit from the status quo order—i.e., the U.S.—would exert more effort to stop their production (Lake, 1993). Other countries have been more brazen with their assistance decisions (Kroenig, 2009).

gathering operation. The first transfers came in the form of radioisotope production following the Atomic Energy Commission's establishment in 1946. Prior to that, international scientists would obtain radioisotopes from U.S. research universities. The Atomic Energy Act shut that down. AEC Chairman Lewis Strauss voiced the central concern that scientists abroad could use those shipments to improve their ability to construct nuclear weapons, at the detriment to U.S. national security (Krige, 2006, 169). The hardline position changed when policymakers concluded that such donations could "open doors" for intelligence gathering and "loosen tongues" (Krige, 2006, 172).

This realization came at a critical time. *Science and Foreign Relations*, a 1950 State Department memo, described intelligence estimates of scientific capabilities as "woefully inadequate" and recommended that the U.S. improve its awareness of foreign scientific developments.[5] The report contained a secret supplement detailing the solution: "employing civilian and diplomatic channels to obtain scientific intelligence intended to aid national security planning," thereby enrolling American scientists into espionage programs (Doel, 1997, 219). For them, "scientific internationalism and scientific intelligence were two sides of the same coin" (Krige, 2006, 167). American scientists were to develop trust with their international colleagues and then exploit "that trust to learn what others were doing, to establish the limits of what they could speak about freely, and to assess the dangers that may lurk behind what they left unsaid" (Krige, 2006, 167).

This was the start of what became a broader theme. By the end of that year, the United States had established the Office of Scientific Intelligence within the CIA and similar organizations within the State Department and Department of Defense (Doel, 1997, 217). The Second Hoover Commission would later echo the recommendation and push for the CIA to deploy attachés to collect scientific intelligence (Lexow, 1966, 23).

Under the scope of these calls, President Eisenhower implemented the Atoms for Peace program with scientific intelligence as one of its major goals. The first Atoms for Peace conference provides an illustration of the intelligence benefits. Held in Geneva in 1955, it represented a unique opportunity. More than a thousand delegates attended. While American officials used the platform to showcase civilian nuclear research, the intelligence community operated backstage. The conference "lifted the veil of secrecy from the reactor programs in the industrialized countries" and served as a source of

[5] See Needell 2013 (141-149) for background on the memo.

"scientific intelligence gathering" (Krige, 2006, 166). Scientists had the opportunity to "probe into the laboratory life of others, ... assess the quality of what [other scientists] were doing" (Krige, 2006, 166), and "judge their competencies" (178). This was one of many scientific intelligence programs that the United States and other countries initiated during the Cold War (Doel, 2010).

The assistance-as-intelligence strategy expanded over time. A post-mortem report (CIA, 1974*a*) on India's nuclear test provides some background. It concluded that acquiring actionable intelligence "should be given much higher priority than it seems to have enjoyed to date." Furthermore, echoing the recommendation to exploit scientists' personal contacts, the report described human intelligence as "clearly the key" in correctly anticipating the next nuclear threats. Technical specialists were the missing link in the process. "At a minimum," it recommended that "such specialists periodically visit *nth* countries" to "backstop the community's collection efforts and ensure in particular that [human intelligence] collectors understand the relationship of their role" to technical systems.[6]

2.3 Scientific Intelligence Gathering in Action

Although it is difficult to ascertain the exact purpose for any given assistance program, bilateral exchanges were regularly a front for human intelligence gathering and the aforementioned specialist visits. The United States' approach to Taiwan provides a helpful example.[7] U.S. intelligence had deduced broad aspects of Taiwan's nuclear activities but wanted a better sense of the severity of the threat. In a meeting with the head of Taiwan's Atomic Energy Council, the American mission in Taipei proposed that the U.S. send a delegation of top nuclear scientists to the island with an official purpose of strengthening civilian nuclear cooperation between the two sides.[8] However, a diplomatic cable reveals a more pressing purpose (State Department, 1973*b*). "In addition to the ostensible purpose cited above," it states, the "team would have further

[6] "Nth" country is shorthand for the hypothetical next country to develop nuclear weapons.

[7] This effort with Taiwan and some of the others below pre-date the post-mortem report but were the basis of what the CIA wanted to expand upon.

[8] In cases like this, the intelligence gathering portion of the mission did not result in a meaningful change in the recipient's capacity. Nevertheless, the technology sharing relationship facilitated the mission, and thus it fits within the scope of my argument. Regardless, a situation where no transfer occurs is a special case of the model I develop below.

objective of acquiring information about identity and progress of ROC coterie which advocates development of nuclear weapons capability ... It would seek to talk to selected persons knowledgeable about ROC activities in this area and to visit all sites of interest to us." As with many other intelligence gathering missions, the U.S. relied on "well-qualified people from the academic world" to handle the mission (State Department, 1973*a*).

The new information acquired help shape the 1974 "Prospects for Further Proliferation of Nuclear Weapons" Special National Intelligence Memo (Richelson, 2007, 270). It concluded that Taiwan had the capacity to develop a small weapons program. However, the U.S. could determine Taiwan's next steps because "support for the island's security and attitudes about the possibility of a nuclear-armed Taiwan" were paramount. Thus, the new information helped guide the United States on how it should negotiate with Taiwan.

Other bilateral examples abound. The United States debriefs American businesses as they work on foreign nuclear facilities. For example, according to a 1975 diplomatic cable from Seoul, the General Electric team sent to South Korea was "impressed with knowledge of nuclear matters displayed by [South Korean] officials they met with" (State Department, 1975). Of particular concern was that "[plutonium] reprocessing was out of idea stage and now a definite plan." The mounting evidence of South Korea's competency forced the intelligence community to reconsider the country as a potential proliferator; only six months earlier, the Special National Intelligence Estimate had treated South Korea as an afterthought (CIA, 1974*b*).

For Israel, President Kennedy tapped Abraham Friedman, the science attaché at the U.S. embassy in Israel. His exchanges with Israeli scientists indicated that Israel had command over nuclear technology. However, there were also indications of "serious deficiencies" that needed to be remedied before the country could "participate in either civilian or military applications of atomic energy" (State Department, 1963). A University of Michigan professor with access to Ernst David Bergmann, father of the Israeli bomb, also provided useful information (Cohen, 1998, 87).

For India, President Johnson dispatched Jerome Wiesner, the former head of the Kennedy administration's science advisory committee. Wiesner's goal was to determine India's proliferation capacity at a time when Homi Bhabha, chairman of India's Atomic Energy Commission, was making bold claims about his country's readiness

(Sarkar, 2015, 937-938; Krige and Sarkar, 2018, 258-259). Over the long term, the information the U.S. received accurately captured India's ability to build nuclear weapons (Richelson, 2007, 228–229).

Although the availability of U.S. historical documents have tilted the conversation in that direction, assistance-as-intelligence is not unique to the United States. Norwegian and French scientists also provided information on Yugoslavia's nuclear cadre (Hymans, 2012, 177–178, 182). Canada's work with India provided information about the latter's willingness to proliferate (State Department, 1972*b*). Beyond that, in perhaps the most formalized exchange, Argentina and Brazil developed a joint commission in the 1980s. They also have provided assistance to each other using their relative strengths, with Argentina supplying uranium concentrate and Brazil producing pressure vessels (Doyle, 2008).

Thus far, the framing of assistance-as-intelligence may give the impression that this information gathering only leads to better decisions. Unfortunately, intelligence is noisy, couched in terms of what experts observed and what they could extrapolate from that. Sometimes this speculation is correct. Other times, the information is misleading and causes officials to formulate the wrong policy in response.

The United States' experience with Israel in the late 1950s provides an example. In short, the United States underappreciated Israeli capacity and did not put much effort into reaching an agreement. The "Post-Mortem" of Special National Intelligence Estimate 100-8-60 revealed where U.S. intelligence went wrong (Cohen, 1998, 81–85). One primary source of information for the United States was a U.S.-Israeli bilateral site. The signs pointed there that Israel had a baseline level of competency but that it was not a major threat to break the proliferation barrier. Despite other intelligence that suggested that Israel might be better prepared to proliferate, the U.S. trusted the information it had coming through its assistance program. "The general feeling [was] that Israel could not achieve" the capabilities that the other source suggested "without outside aid from the US or its allies" and discounted rumors of French assistance (Cohen, 1998, 84). French assistance was forthcoming, however, and the poor information caused the United States to miss its best chance at stopping Israel's bomb (Cohen, 1998, 79–98). As such, intelligence from nuclear assistance proved counterproductive and played a role in Israel's development of nuclear weapons.

2.4 Comparison to Existing Theories

Wrapping up, I incorporate all of these incentives—bargaining, technical assistance, and noisy signals—into my model. The results indicate that information acquisition is sometimes worth giving a would-be proliferator a faster and cheaper route to nuclear weapons.

In comparison to existing explanations, mine has three 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.[9] The net utility gain resolves that puzzle. Second, it helps explain why assistance has persisted into the post-Cold War era despite a lack of superpower competition.

Finally, many studies find pathways by which assistance can reduce proliferation outcomes. Hymans (2012) motivates how assistance programs causes brain drain, ultimately *decreasing* a state's competency. Kemp (2014) describes how sellers may give inferior technology, leaving the recipient with a worse outcome than had it created a native project. Montgomery (2013) argues that assistance reduces proliferation risks when on the recipient was already pursuing a bomb. These are sensible arguments, but they do not match Fuhrmann's (2009*a*) broader statistical connection between assistance and proliferation activities.[10] In contrast, my mechanism can counterintuitively predict more proliferation following assistance despite the information gain. The theoretical results also indicate that even if the net effect of providing assistance to pursuers reduces proliferation (Montgomery, 2013), parameters exist where this is not the case. In turn, this should give the nonproliferation regime pause about when it should encourage donor states to provide assistance.

To be clear, my model only covers cases where bargaining and intelligence incentives are at play. The above discussion suggests that a significant number of critical cases fall into this category. This is the type of tradeoff present in all models, but it is worth

[9]Fuhrmann (2009*a*) suggests that the purposes was to gain influence within the recipient country, but it is still unclear why a state would gain influence through assistance and not a different manner that lacks the same drawback.

[10]Bluth et al. (2010) provide an alternative account: engaging in proliferation behaviors causes states to receive more assistance agreements, as new deals replace the broken ones. Fuhrmann (2012, 19) notes that the percentage of canceled agreements is small and that those canceled still usually result in some transfer. My paper does not adjudicate the statistical issue. However, it does provide reassurance that Fuhrmann's result could arise with rational policymakers.

reiterating that the model does not imply that all assistance decisions look like the structure I analyze.

3 The Model

The game consists of two states, a **D**onor and **R**ecipient, negotiating over the latter's decision to build nuclear weapons. Nature begins by drawing R as a "low" type L with probability p and a "high" type H with probability $1 - p$. R observes the realization but D does not. This is the source of incomplete information, and all other parameters are common knowledge.

D then chooses whether to provide assistance. If it does, Nature sends D a noisy signal about R's type. In particular, Nature reports 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.[11]

Regardless of D's assistance decision and the message possibly received, a round of bargaining follows. Here, D offers an ultimatum $x \geq 0$. This captures the total package of transfers R receives if and only if it forgoes nuclear weapons pursuit.[12] Consistent with recent work on nuclear negotiations (Debs and Monteiro, 2014; Miller, 2014; Spaniel, 2019; Mehta, 2020), this includes improved diplomatic relations, sanctions relief, alliance benefits, military assistance, and economic aid.[13] R can accept or reject. The game ends either way.

Payoffs are as follows. Regardless of how negotiations unfold, D pays $k \geq 0$ if it provides assistance. This parameter captures the idea that assistance served many purposes. For example, if superpower competition gave a propaganda benefit to D for providing assistance, then this would implicitly make k lower. In fact, the key

[11]One might imagine that the strength of the signal depends on the type realization—e.g., Nature almost always correctly reveals a low type but conveys noise for a high type. This would not alter the key substantive takeaways, but it would change the calculations of the equilibrium cutpoints.

[12]The appendix formalizes negotiations if R could enjoy the concessions while secretly building a weapon and then enjoy the benefits of proliferation thereafter. Consistent with previous work (Spaniel, 2019, 126–156), agreements still work if proliferation is sufficiently costly, though D must pay a premium. D also faces the risk-return tradeoff I describe below, where it makes risky offers if its belief that R is the low type is sufficiently large. This generates the same key implication of the baseline model: transfers may or may increase proliferation rates, depending on D's prior belief.

[13]In fact, the United States codified this idea with Arms Export Control Act, which prohibits such transfers to states engaging in unsafeguarded fissile material production.

theoretical results go through even if the net "cost" were negative, meaning that the informational benefit is a secondary concern to D.[14]

Beyond the k payoff component, 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 assistance it may have received. If R rejects, however, both of those points come into play. R's proliferation speed and cost are a function of its type and the transfer. Let $\delta_i(\bullet) \to [0,1]$ map an underlying skill and transfer (or lack thereof) for a given type $i \in \{L, H\}$ to a delay. I assume that the function is strictly increasing. Denote the possible transfer as $t \geq 0$, the low type's skill as $s > 0$, and the high type's as $s' \geq s$.[15] Then I write the delay for a high type without assistance as $\delta_H(s')$ and with assistance as $\delta_H(s'+t)$. The low type's delays are analogous. To ensure that the high type develops weapons at least as fast as the low type, let $\delta_H(\bullet) \geq \delta_L(\bullet)$.[16]

In addition, $c_i(\bullet) \to (0, \infty)$ maps an underlying skill and transfer (or lack thereof) for a given type $i \in \{L, H\}$ to a cost of proliferation. I assume that the function is strictly decreasing. Like before, the high type's without assistance is $c_H(s')$, the high type's with assistance is $c_H(s' + t)$, and so forth.[17] As with other models of conflict, this term implicitly captures R's willingness or resolve to develop nuclear weapons. To capture that, $c_H(\bullet) \leq c_L(\bullet)$—that is, for the same total level of competency, the high type's realized cost is weakly less than the low type's.

Let $b > 0$ represent the coercive value, prestige gain, or domestic benefit from obtaining nuclear weapons. Then R's payoff equals $\delta_i(\bullet)b - c_i(\bullet)$, 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 ne-

[14]Intuitively, D would provide assistance under a wider range of circumstances; how that impacts the bargaining outcome depends on D's prior belief. I keep k positive because it reduces the number of cases I must sort through.

[15]As mentioned above, the case where $t = 0$ represents cases where the mission provides no direct benefit to R but gives D a noisy signal. This would better connect to Montgomery's (2013) expectation that assistance does not much help with the proliferation process or a world where the Nuclear Suppliers Group's purpose is completely fulfilled; higher values instead connect with Fuhrmann's (2012) main claims. Meanwhile, $s' = s$ represents the case where the types have the same skill but differing resolve.

[16]For cases where a type's resolve has no impact on delay, $\delta_H = \delta_L$ for all identical inputs. I allow for the inequality to be weak because whether competence and resolve are correlated are inconsequential for the propositions' claims.

[17]For cases where the types have identical resolve, $c_H = c_L$ for all identical inputs

gotiations between friends, and environmental damage or risk of catastrophic accident. Then D's payoff equals $-\delta_i(\bullet)e$.

Recapping, the game proceeds as follows:

1. Nature privately informs R of its type
2. D chooses whether to provide assistance
3. If D provided assistance, Nature sends it a noisy signal, indicating either "low" or "high"
4. D offers concessions to R
5. R accepts those concessions or proliferates, with proliferation cheaper and faster if D provided assistance

Before continuing, a few notes about the game's assumptions are in order. First, under this setup, R automatically receives the transfer if D offers 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 low type would have incentive to reject to avoid revealing its lack of skill. However, transfers exhibit an "unraveling" principle. The high 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 low type declines, it reveals its type by virtue of *not* accepting assistance. This unraveling effect also applies to distributions beyond the 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 donors 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 existing many models (Debs and Monteiro, 2014; Bas and Coe, 2016; Spaniel, 2019). Preventive war is a critical feature of some interactions, but it is not universal. When the U.S. bargains with allies, for example, policymakers in Washington would rather permit these countries to develop a nuclear weapon than fight a war and ruin the friendship.[18] Sworn enemies may also decline preventive action when the target has a large conventional deterrent or the protection of an 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, which is a key property of "experimental" models (Paine and Tyson, 2019).

To be clear, the lack of a credible threat for preventive war is a scope condition of my argument. Indeed, if preventive war is borderline-credible, would-be proliferators have an incentive to *underrepresent* their proficiency. This is because preventive war looks more attractive to declining states when the shift occurs more rapidly, as they internalize the negative repercussions sooner. Consistent with this, American assistance programs saw more early successes with allies (i.e., those not at risk of preventive war) than with antagonists (Lexow, 1966).

However, the incentives flip when preventive war is incredible, and examples of states wanting to reveal capacity are plentiful. In the early 1990s, for instance, Pakistan became desperate to convince the United States that it could easily put together a weapon if it wished to (Ahmed, 1999, 190–191). North Korea took similar steps in the early 2000s with their revelations to Siegfried Hecker (Hecker, Braun and Lawrence, 2006). Taiwan even offered to pay for American experts to take residence on the island (State Department, 1976). These sorts of behaviors only make sense if the recipient does not worry about preventive war, and I focus on such cases.

For the remainder of the paper, I also assume that each type has a non-negative value for building nuclear weapons in the absence of a transfer.[19] 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 couple of exceptions. First, 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 low type to not proliferate without offering any concessions, making the functional cost

[18] Alliance patrons also punish protégés for proliferating, which creates disutility for both parties (Bas and Coe, 2018). This most often takes the form of sanctions. To address that, the appendix develops an extension with sanctions. It shows that high proficiency states still want to reveal information in the manner described. The intuition is straightforward here. If a state could quickly construct a weapon, then there is not enough time for sanctions to have an effect.

[19] Formally, this requires $\delta_L(s)b - c_L(s) \geq 0$.

of acquiring the signal lower. As such, focusing on the interior solution also stacks the deck against a transfer occurring. A second 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. This has important substantive implications, which I return to later.

In addition, I 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.[20] This automatically holds if states are unitary actor security rivals and whatever benefits R receives after acquiring weapons implies an equivalent security loss to D. It also applies to allies when D 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 South Korea—a close ally and technically proficient country—if policymakers in Washington did not care about these externalities.

4 Solving the Game

Now to solve the game. Perfect Bayesian equilibrium is the appropriate solution concept. The interaction hinges on whether D performs better by providing assistance. I do this in three parts. The first discusses some of the general incentives that D faces. The second solves for the game when its prior belief places great weight on R being the low type. The final part solves for the game when its prior belief places great weight on R being the high type.

4.1 When Should D Provide Assistance?

To begin, if D does not provide assistance, it faces a classic risk-return tradeoff. It chooses from one of two options. First, it can offer a generous settlement designed to induce both types to accept. This guarantees that D avoids suffering the externality, but it also means that D overpays the low type for its compliance. Alternatively, D

[20]Formally, this requires D's minimum payment to the high type following a transfer to be less than the time-adjusted externality it suffers, or $\delta_H(s'+t)b - c_H(s'+t) < \delta_H(s'+t)e$.

could offer a stingy amount designed to induce only the low type to accept. This implies giving fewer concessions but also results in D suffering the externality whenever R is the high type.

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

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

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

Thus, D's prior belief determines whether proliferation occurs. If it is sufficiently skeptical of R's ability and desire, it makes a smaller offer; if it is instead sufficiently worried, it makes a larger offer. D's prior belief has downstream consequences on whether it wants to provide assistance and the empirical implications that come with the transfer.

Nevertheless, similar 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 proficiency means that D must make greater concessions to buy off either type. For example, the high type's payoff for proliferating moves from $\delta_H(s')b - c_H(s')$ before assistance to $\delta_H(s'+t)b - c_H(s'+t)$ after. Because development speeds increase and costs decrease in transfers, the latter payoff is larger. Therefore, D must increase its offer to induce the high type to accept. The same holds for the low type.

Assistance also forces D to update its beliefs. Intuitively, receiving a high signal increases D's belief that it is facing a high 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 low type.[21] All else equal, this makes D more inclined to make the risky offer.

From here, D calculates its expected payoff for giving assistance and compares it to its payoff for making no transfer. This is computationally intensive, so I leave that

[21]See Lemma 6 in the appendix for proof.

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 assistance is that the sets of types D induces to accept with its offers with and without the signal are not identical.*

Put differently, if D provides assistance, one of the signals must be persuasive enough to change the type D targets.[22] For example, without assistance, imagine that D would make the risky offer that only the low type accepts. However, upon receiving the high signal, D would instead make the safe offer.[23] Then the signal is persuasive because it swayed D's belief enough that the offer no longer targets the same type of R. If D would continue offering an amount that just the low 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 had originally planned to appease. As such, in equilibrium, D only provides assistance to obtain persuasive signals.

Lemma 3 provides an immediate empirical implication. One might suspect that intelligence estimates of countries who receive assistance would be better than intelligence estimates of countries who do not. However, the model demonstrates a selection effect. The extreme case helps explain why. If D had complete information, it would not provide assistance—doing so improves R's outside option without telling D anything it did not know already. But the same holds for cases where D's prior is close to 0 or 1—that is, close to complete information—because any given signal is unlikely to change how D pursues the post-transfer bargaining. In turn, intelligence problems may be worse for the cases where we observe intelligence gathering.[24]

[22]This result does not survive for sufficiently negative "costs". For instance, if D received an infinite benefit just for providing assistance, it would do so even if its bargaining position did not change.

[23]Note that the "risky" and "safe" offers differ depending on whether R received assistance. I use these phrases here to describe whom D is attempting to appease rather than the size of the offer.

[24]It is true, of course, that a case with assistance gives D better information than in a counterfactual world where D did not give that recipient assistance. It is also possible that assistance can generate better information across cases if the strength of the signal is sufficiently high.

Lemma 4. *(Responsiveness) A necessary condition for D to provide assistance 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.[25] 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 low 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 high type's skill with assistance. However, suppose instead that D did not make a transfer. Then D can still buy off both types. But the amount necessary to do that is a quantity commensurate with the high type's skill *without* assistance. This is cheaper and therefore better.

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 low 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 the high type's speed consistent with with assistance. However, suppose instead D did not make a transfer. It can still buy off the low type, but now the payment need only be commensurate with the low type's skill *without* assistance. Meanwhile, the portion of the time R is high and rejects, D suffers the externality at the high type's speed consistent 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 low following the corresponding signal must be higher than its prior. Likewise, that belief following the opposite 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.

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

Proposition 1. *Suppose the signal is sufficiently strong compared to the cost of assistance. Then D provides assistance.*

[25] Like before, this does not survive sufficiently negative "costs". For example, even if each signal led D to the same demand, an infinite benefit would render the informational reward moot.

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. Assistance 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 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. One can see by observing that as the signal becomes perfectly informative, D's belief that R is the low type 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 the cutpoint determining whether D offer preference.

If the signal is sufficiently strong, then D compares its payoffs making and not making the transfer.[26] 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 huge concessions. The direct monetary cost of the transfer can also make assistance suboptimal.

4.2 When D Is Skeptical

Overall, Proposition 1 explains why donors 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 therefore cut the discussion in two, beginning with cases when D is skeptical:

Proposition 2. *Suppose D begins skeptical (i.e., $p > p^*$). Then the probability of proliferation is strictly lower when D provides assistance. Moreover, the probability of*

[26] This step is the same regardless of whether D's prior belief makes it skeptical or worried, though the actual calculations are different.

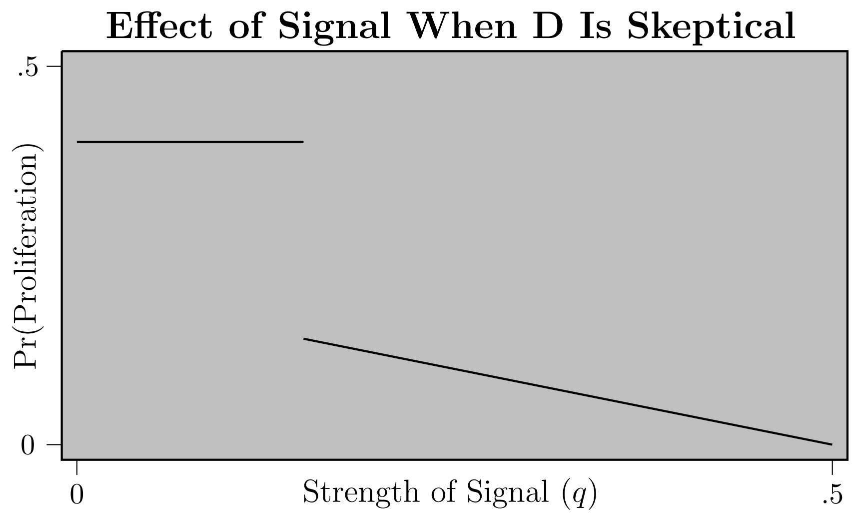

Figure 1: The probability of proliferation when D begins as skeptical.

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 high type would reject. However, if D provides assistance in equilibrium, a high signal must induce D to appease the high type. Thus, the only time R proliferates is when D wrongly receives the low signal and R is actually high. 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 low type, causing the high 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 high.

Nevertheless, further increases to the signal quality make the transfer worthwhile.[27] At that point, the probability of proliferation drops discontinuously because D finally

[27] 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.

releases all the pent-up signal value.[28] Two things must now go wrong for proliferation to occur—D must initially suspect the wrong opposing type *and* the signal must not correct D's inaccurate belief. Before, only D's initial suspicion had to be wrong for proliferation to occur. Moreover, as the signal becomes stronger, the wrong signal becomes less likely. 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.

To substantively ground these ideas, consider how the United States handled Indonesia's nuclear threats.[29] In the mid-1960s, Indonesia was not on the United States' nuclear radar. The default assumption was an inability to construct nuclear weapons. Nevertheless, following China's 1964 test, Brigadier General Hartono publicly declared that Indonesia was only a year away from developing a bomb. Taken aback, the United States began using its scientific intelligence assets to reassess the situation. This was possible because the United States had installed Indonesia's only reactor as a part of the Atoms for Peace initiative, a TRIGA-Mark II model.

The information that came back painted a clear picture: the default assumption of low ability to construct nuclear weapons was accurate. The reactor was not suitable for making a nuclear weapon: it had a maximum of six kilograms of 20% enriched uranium, short on both the mass and concentration of U-235 necessary for a single bomb. Indonesia was also lacking a reprocessing facility to extract plutonium from the spent uranium. Human intelligence also painted a clear picture of the country's technical knowledge. Despite only having one reactor, Indonesia struggled to staff it with a competent crew. The head of the program was a medical doctor, not a proper nuclear engineer, and there were perhaps only a dozen Indonesians in total with real training. What little training occurred within the country focused on nuclear-adjacent positions rather than weapons building. Moreover, the general "[didn't] know what he was talking about" when it came to building nuclear weapons, and that for Indonesia to acquire a nuclear weapon, "it would have to be given to [them]."

Indonesia continued blustering about nuclear weapons, with President Sukarno making empty promises starting in July 1965. The United States, with confirmation from its intelligence sources, did not offer serious concessions to stop Indonesia, believing

[28]If no signal is worthwhile, then the probability of proliferation remains flat across the entire range.
[29]The narrative I describe below is a summary of Cornejo (2000, 33–39).

that they were unnecessary. By September, Sukarno halted his proclamations. An aborted coup in October swept Sukarno out of power and ended the Indonesian nuclear arc.

4.3 When D Is Worried

The previous case makes it appears that assistance has a chilling effect on proliferation. But another parameter space exists where D's default inclination is to make safe proposals. The following proposition shows that the chill is not universal:

Proposition 3. *Suppose D begins worried (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 worried, its optimal offer guarantees a nonproliferation outcome. As such, any manipulation to the bargaining environment can only promote proliferation. Here, this happens because the signal may be misleading. A low signal nudges D to make an aggressive offer that the high type would reject. And while D correctly updates its belief that the low type is more likely than it thought previously, the signal is not perfect. Whatever portion of the time the signal inaccurately indicates the low type, 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.

However, increasing the signal strength beyond a certain point causes D to transfer assistance.[30] This causes a discontinuous *increase* in the probability of proliferation. Specifically, the right side of the graph depicts the probability R is actually the high type but D falsely receives the low signal. D then tailors its offer to the wrong type, leading the high type to reject. The probability of proliferation decreases as the signal

[30]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.

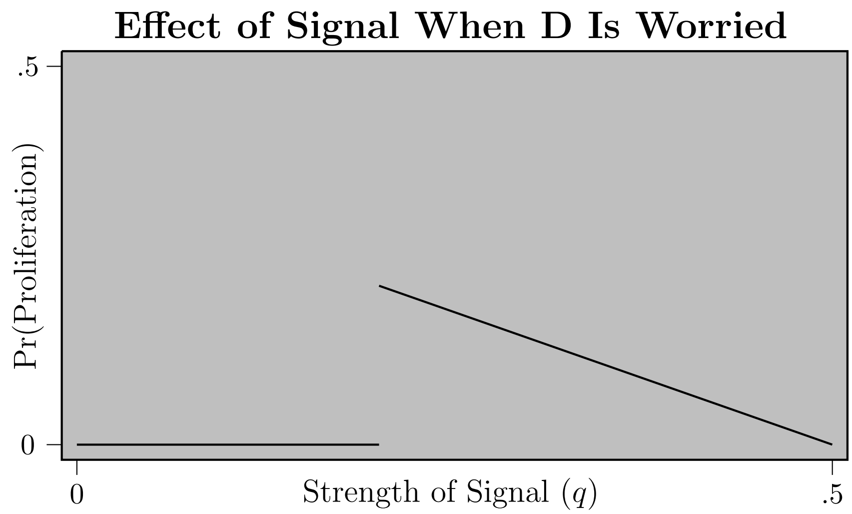

Figure 2: The probability of proliferation when D begins as worried.

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.

Zooming out, although Proposition 3 produces a similar empirical implication as the "spreading temptation" logic, my model's mechanism is distinct. Fuhrmann (2009*a*) 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 offer washes out lower costs and faster development speeds. Instead, the signal can cause D to become more optimistic that R is the low type, and that belief leads D to pursue riskier offer strategies.

Diplomacy in the lead up to India's nuclear test reflects this logic. Of all the cases described here thus far, U.S. intelligence efforts had perhaps best captured India's technical capacity. There was "no question that India could proceed rapidly and with little difficulty" (State Department, 1972*a*). Western sources had plenty of contact with Indian nuclear scientists and were aware of the key facilities that would contribute to the test. Thus far, the intelligence was pointing the United States in the right direction.

However, the failure came in assessing India's willingness to proliferate. Here, the estimates missed the mark, with human intelligence from civilian nuclear sources contributing to the problem. In 1972, James Lorne Gray spoke with Homi Nusserawnji

Sethna, in a meeting between the chairmen of the Canadian and Indian nuclear energy departments, respectively. Gray's takeaway from the conversation was that "there was no chance of India launching into a device program as long as Sethna and Ghandi remain[ed] in office," referring to India's prime minister at the time (State Department, 1972*b*).

Canada subsequently did not take serious active steps to further convince India not to proliferate. And although U.S. intelligence worried there was some chance that Gray had misread Sethna (State Department, 1972*c*), no firm intelligence came in to contradict the notion. As a result, Washington also adopted a casual approach. By 1974, the estimated chances of proliferation had *declined* (State Department, 1974). Yet four months later, India detonated the Smiling Buddha, leaving the U.S. intelligence community shocked. In retrospect, a major problem was the small footprint of the program. Only 75 people worked to construct the weapon (Perkovich, 2001, 172). Thus, there were fewer opportunities for the U.S. to receive accurate information and more chances for ignorant Indian scientists to give false indication of disinterest in weaponization. Correspondingly, the post-mortem suggested that the U.S. double down on human intelligence (CIA, 1974*a*).

Taking stock, the results depicted in Figures 1 and 2 indicate that the relationship between assistance and proliferation can break either way. Nevertheless, there are two reasons to think that the empirical record would show a positive correlation. First, this discussion has focused on the interior solution, where R's native competency is high enough that it finds proliferation profitable. However, when both types have negative proliferation utilities, D makes no transfer, and R never proliferates. This creates a broad empirical baseline of countries that both do not receive assistance and do not develop nuclear weapons.

Second, there is theoretical reason to think that Figure 2 covers a substantial portion of the remaining empirical cases. Consider the cutpoint that determines which parameter space the game falls in: $p^* = \frac{\delta_H(s')e - \delta_H(s')b + c_H(s')}{\delta_H(s')e - \delta_L(s)b + c_L(s)}$. 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 issues point in that direction. For one, as e increases, p^* tends toward 1. Substantively, if the externality of proliferation is large, D must be confident R is the low type to make a risky offer. The substantive literature suggests that this is often the case. The problems here range from nuclear accidents and the possibility of catastrophic inadvertent war (Sagan, 1995), to loose nuclear materials, to emboldened alliance partners, to environmental externalities, to zero-sum security losses. 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.

For another, p^* also increases as b decreases. Substantively, if the benefit of proliferation is small, the difference in reservation values merge together. Under this situation, the premium D must pay to the low type to guarantee that the high type accepts is small. Consistent with other analysis of this type of premium (Reed, 2003; Arena, 2013; Spaniel and Malone, 2019), D must be very sure R is the low type to make the risky offer that might cause it to suffer the externality. This 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. Consistent with this, Fuhrmann (2009*a*) finds that assistance correlates with more proliferation.

5 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 transfers force opponents to make deeper concessions later, the price is sometimes worth the benefit.

The theory helps explain a curious empirical finding in the proliferation literature. States receiving nuclear cooperation agreements from their peers 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 expected welfare.

The source of the selection problem indicates a major challenge in teasing it out of the data. For example, one might wish to operationalize belief through capacity variable, with lower capacity representing skepticism. This is *not* valid, however. Measures like Jo and Gartzke (2007) and Smith and Spaniel (2018) generate a mean expected proficiency, whereas the model requires a belief over a range of possible beliefs. Although quantitative scholars have begun incorporating uncertainty into their regressions (Reed, 2003; Rider, 2013; Spaniel and Smith, 2015), existing proxies only capture the general informational environment. The selection problem here requires defining more specific beliefs—i.e., pessimistic or optimistic—for the uncertain actor. In the absence of an adequate proxy, and given the selection problem at hand, it is not surprising that empirical findings would be mixed. Model specification matters more than the current literature appreciates.

Future research could consider the variation in donor incentives. The IAEA, for example, provides assistance through its technical cooperation program (Brown and Kaplow, 2014). Ostensibly, the IAEA's goal is to reduce proliferation rates, but the organization does not grant concessions like rival states or allies would. This gives the IAEA more incentive to provide assistance—but only when proliferation rates would be high in the absence of new information. As a result, an even greater barrier to inference would appear in the data. With assistance only coming when the probability of proliferation is positive—and still staying positive due to the noisy signal—it would appear that assistance begets proliferation. Yet the same sort of selection problem would suggest otherwise, as cases without a risk of proliferation are predisposed not to receive assistance at all.

References

Ahmed, Samina. 1999. "Pakistan's nuclear weapons program: Turning points and nuclear choices." *International Security* 23(4):178–204. 15

Arena, Philip. 2013. "Costly signaling, resolve, and martial effectiveness." Manuscript, University of Buffalo. 26

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. 2, 4, 14

Bas, Muhammet A and Andrew J Coe. 2018. "Give peace a (Second) chance: A theory of nonproliferation deals." *International Studies Quarterly* 62(3):606–617. 15

Benson, Brett V. 2012. *Constructing international security: alliances, deterrence, and moral hazard*. New York: Cambridge University Press. 5

Bluth, Christoph, Matthew Kroenig, Rensselaer Lee, William C Sailor and Matthew Fuhrmann. 2010. "Civilian nuclear cooperation and the proliferation of nuclear weapons." *International Security* 35(1):184–200. 11

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. 2, 11, 27

CIA. 1974*a*. "Post-mortem report: An examination of the intelligence community's performance before the Indian nuclear test of May 1974." Director of Central Intelligence: `https://nsarchive2.gwu.edu/nukevault/ebb367/docs/7-74%20post-mortem.pdf`. 8, 25

CIA. 1974*b*. "Prospects for further proliferation of nuclear weapons." Director of Central Intelligence: `https://nsarchive2.gwu.edu//NSAEBB/NSAEBB240/snie.pdf`. 9

Cohen, Avner. 1998. *Israel and the Bomb*. New York: Columbia University Press. 9, 10

Colgan, Jeff D and Nicholas L Miller. 2019. "Rival hierarchies and the origins of nuclear technology sharing." *International Studies Quarterly* 63(2):310–321. 1

Cornejo, Robert M. 2000. "When Sukarno sought the bomb: Indonesian nuclear aspirations in the mid-1960s." *The Nonproliferation Review* 7(2):31–43. 22

Debs, Alexandre and Nuno P Monteiro. 2014. "Known unknowns: Power shifts, uncertainty, and war." *International Organization* 68(1):1–31. 2, 12, 14

Debs, Alexandre and Nuno P Monteiro. 2016. *Nuclear Politics*. New York: Cambridge University Press. 4, 15

Doel, Ronald E. 1997. Scientists as Policymakers, Advisors, and Intelligence Agents: Linking Contemporary Diplomatic History with the History of Contemporary Science. In *The Historiography of Contemporary Science and Technology*, ed. Thomas Söderqvist. Paris: Harwood pp. 215–244. 1, 6, 7

Doel, Ronald E. 2010. "Does scientific intelligence matter?" *Centaurus* 52(4):311–322. 8

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. 10

Drezner, Daniel W. 1999. *The sanctions paradox: Economic statecraft and international relations*. New York: Cambridge University Press. 6

Fuhrmann, Matthew. 2009*a*. "Spreading temptation: Proliferation and peaceful nuclear cooperation agreements." *International Security* 34(1):7–41. 1, 2, 4, 6, 11, 24, 26

Fuhrmann, Matthew. 2009*b*. "Taking a walk on the supply side: The determinants of civilian nuclear cooperation." *Journal of Conflict Resolution* 53(2):181–208. 3

Fuhrmann, Matthew. 2012. *Atomic assistance: How Atoms for Peace programs cause nuclear insecurity*. Ithaca: Cornell University Press. 11, 13

Hecker, Siegfried S., Chaim Braun and Chris Lawrence. 2006. "North Korea's Stockpiles of Fissile Material." *Korea Observer* 47(4):721–749. 15

Hersman, Rebecca KC and Robert Peters. 2006. "Nuclear U-turns: Learning from South Korean and Taiwanese rollback." *Nonproliferation Review* 13(3):539–553. 1

Hymans, Jacques EC. 2012. *Achieving nuclear ambitions: scientists, politicians, and proliferation*. Cambridge University Press. 10, 11

Jo, Dong-Joon and Erik Gartzke. 2007. "Determinants of nuclear weapons proliferation." *Journal of Conflict Resolution* 51(1):167–194. 4, 27

Kemp, R Scott. 2014. "The nonproliferation emperor has no clothes: The gas centrifuge, supply-side controls, and the future of nuclear proliferation." *International Security* 38(4):39–78. 11

Krige, John. 2006. “Atoms for Peace, scientific internationalism, and scientific intelligence.” *Osiris* 21(1):161–181. 1, 7, 8

Krige, John and Jayita Sarkar. 2018. “US technological collaboration for nonproliferation: Key evidence from the Cold War.” *The Nonproliferation Review* 25(3-4):249–262. 10

Kroenig, Matthew. 2009. “Importing the bomb: Sensitive nuclear assistance and nuclear proliferation.” *Journal of Conflict Resolution* 53(2):161–180. 3, 6

Kydd, Andrew H. 2010. “Rationalist approaches to conflict prevention and resolution.” *Annual Review of Political Science* 13:101–121. 2

Lake, David A. 1993. “Leadership, hegemony, and the international economy: Naked emperor or tattered monarch with potential?” *International Studies Quarterly* 37(4):459–489. 6

Lexow, Wilton. 1966. The Science Attache Program. In *Studies in Intelligence*, ed. Central Intelligence Agency. Washington, DC: Central Intelligence Agency pp. 22–27. 6, 7, 15

Marinov, Nikolay. 2005. “Do economic sanctions destabilize country leaders?” *American Journal of Political Science* 49(3):564–576. 42

Medhurst, Martin J. 1997. “Atoms for peace and nuclear hegemony: The rhetorical structure of a Cold War campaign.” *Armed Forces & Society* 23(4):571–593. 1

Mehta, Rupal N. 2020. *Delaying Doomsday: The Politics of Nuclear Reversal*. New York: Oxford University Press, USA. 4, 12

Miller, Nicholas L. 2014. “The secret success of nonproliferation sanctions.” *International Organization* 68(4):913–944. 5, 12

Miller, Nicholas L. 2017. “Why nuclear energy programs rarely lead to proliferation.” *International Security* 42(2):40–77. 2

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. 5

Montgomery, Alexander H. 2013. Stop helping me: When nuclear assistance impedes nuclear programs. In *The nuclear renaissance and international security*, ed. Adam Stulerg and Matthew Fuhrmann. Stanford: Stanford University Press pp. 177–202. 11, 13

Montgomery, Alexander H and Adam Mount. 2014. "Misestimation: Explaining US Failures to Predict Nuclear Weapons Programs." *Intelligence and National Security* 29(3):357–386. 2, 4

Needell, Allan A. 2013. *Science, Cold War and the American State*. New York: Routledge. 7

Paine, Jack and Scott A Tyson. 2019. "Uses and abuses of formal models in political science." Manuscript, University of Rochester. 15

Paul, TV. 2000. *Power versus prudence: Why nations forgo nuclear weapons*. Montreal: McGill-Queen's University Press. 4

Perkovich, George. 2001. *India's nuclear bomb: The impact on global proliferation*. University of California Press. 25

Powell, Robert. 1999. *In the shadow of power: States and strategies in international politics*. Princeton: Princeton University Press. 5

Reed, William. 2003. "Information, power, and war." *American Political Science Review* 97(4):633–641. 26, 27

Reiss, Mitchell. 1988. *Without the bomb: The politics of nuclear nonproliferation*. New York: Columbia University Press. 4

Richelson, Jeffrey. 2007. *Spying on the bomb: American nuclear intelligence from Nazi Germany to Iran and North Korea*. New York: WW Norton & Company. 9, 10

Rider, Toby J. 2013. "Uncertainty, salient stakes, and the causes of conventional arms races." *International Studies Quarterly* 57(3):580–591. 27

Sagan, Scott Douglas. 1995. *The limits of safety: Organizations, accidents, and nuclear weapons*. Princeton: Princeton University Press. 26

Sarkar, Jayita. 2015. "The Making of a Non-Aligned Nuclear Power: India's Proliferation Drift, 1964–8." *The International History Review* 37(5):933–950. 10

Schwartz, Stephen I. 2011. *Atomic audit: The costs and consequences of US nuclear weapons since 1940*. Washington, DC: Brookings Institution Press. 4

Sechser, Todd S and Matthew Fuhrmann. 2017. *Nuclear Weapons and Coercive Diplomacy*. New York: Cambridge University Press. 26

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 ν-CLEAR: A latent variable approach to measuring nuclear proficiency." *Conflict Management and Peace Science, Forthcoming* . 4, 27

Spaniel, William. 2019. *Bargaining over the bomb: The successes and failures of nuclear negotiations*. New York: Cambridge University Press. 2, 4, 12, 14

Spaniel, William and Bradley C Smith. 2015. "Sanctions, uncertainty, and leader tenure." *International Studies Quarterly* 59(4):735–749. 27, 42

Spaniel, William and Iris Malone. 2019. "The uncertainty trade-off: Reexamining opportunity costs and war." *International Studies Quarterly* 63(4):1025–1034. 26

State Department. 1963. "Israeli program for nuclear reactors." Diplomatic Cable: `https://nsarchive2.gwu.edu//dc.html?doc=5983542-National-Security-Archive-Doc-07-U-S-Embassy`. 9

State Department. 1972*a*. "India to go nuclear?" Bureau of Intelligence and Research Intelligence Note: `https://nsarchive2.gwu.edu/nukevault/ebb367/docs/1-14-72.pdf`. 24

State Department. 1972*b*. "India's nuclear intentions." Diplomatic Cable: `https://nsarchive2.gwu.edu/nukevault/ebb367/docs/3-7-72.pdf`. 10, 25

State Department. 1972*c*. "India's nuclear intentions on South Asia situation." Diplomatic Cable: https://nsarchive2.gwu.edu/nukevault/ebb367/docs/3-14-72.pdf. 25

State Department. 1973*a*. "Reported ROC nuclear weapons development program." Memorandum: https://nsarchive2.gwu.edu/NSAEBB/NSAEBB20/docs/doc11.pdf. 9

State Department. 1973*b*. "ROC nuclear intentions." Diplomatic Cable: https://nsarchive2.gwu.edu/NSAEBB/NSAEBB20/docs/doc30.pdf. 8

State Department. 1974. "India's nuclear intentions." Diplomatic Cable: https://nsarchive2.gwu.edu/nukevault/ebb367/docs/1-19-74.pdf. 25

State Department. 1975. "ROK nuclear program." Diplomatic Cable: https://nsarchive2.gwu.edu//dc.html?doc=3535258-Document-02A-U-S-Embassy-Seoul-Telegram-1089-to. 9

State Department. 1976. "ROC's nuclear intentions: Conversation with Premier." Diplomatic Cable: https://nsarchive2.gwu.edu/nukevault/ebb221/T-9a.pdf. 15

Volpe, Tristan A. 2017. "Atomic leverage: Compellence with nuclear latency." *Security Studies* 26(3):517–544. 4

6 Appendix

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

Lemma 5. *Upon receiving the low signal, D's posterior belief that R is the low type is* $\frac{p(.5+q)}{.5-q+2pq}$*. Upon receiving the high signal, D's posterior belief that R is the high type 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 the low type is p. D can receive a low signal in two ways: R is actually the low type and it received the correct signal with probability $.5+q$, or R is actually the high type and it received the wrong signal with probability $.5-q$. Therefore, the probability of R being the low type upon receiving the low 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 the low type and it received the incorrect signal with probability $.5-q$, or R is actually the high type and it received the correct signal with probability $.5+q$. Therefore, the probability of R being the low type upon receiving the high signal is:

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

□

Lemma 6. *D's posterior belief that R is the low type increases following a low signal and decreases following a high signal (i.e.,* $\underline{r} > p$ *and* $\overline{r} < p$*).*

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

$$\frac{p(.5-q)}{.5+q-2pq} < p$$
$$p > 0$$

This is also true.

$\square$

Most of the remaining sections prove the main results. I also include a brief discussion of an extension in which sanctions impact the proliferator's likelihood of success.

6.1 Proof of Lemmas 1 and 2

Suppose D did not provide assistance. Then the low type accepts if $x \geq \delta_L(s)b - c_L(s)$, and the high type accepts if $x \geq \delta_H(s')b - c_H(s')$.[31] D's optimal offer must therefore be either $\delta_L(s)b - c_L(s)$ or $\delta_H(s')b - c_H(s')$. Any more than $\delta_H(s')b - c_H(s')$ is a needless concession, any less than $\delta_L(s)b - c_L(s)$ 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 low type and rejection from the high type.

We can find D's optimal choice by comparing its utilities for each of these choices. Offering $\delta_H(s')b - c_H(s')$ induces acceptance from both types. D therefore receives the negative of that amount as its payoff. Offering $\delta_L(s)b - c_L(s)$ induces the low type accept but the high 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 high type's pace. Taking the expectation and comparing utilities, the risky offer is better if:

$$-p(\delta_L(s)b - c_L(s)) - (1-p)\delta_H(s')e > -(\delta_H(s')b - c_H(s'))$$
$$p > \frac{\delta_H(s')e - \delta_H(s')b + c_H(s')}{\delta_H(s')e - \delta_L(s)b + c_L(s)} \equiv p^*$$

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

[31] For the usual reasons, no equilibria exist in which a type accepts with any probability less than 1 when indifferent.

6.2 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 $\delta_L(s + t)b - c_L(s + t)$ after receiving the low signal if:

$$\underline{r} > \frac{\delta_H(s' + t)e - \delta_H(s' + t)b + c_H(s' + t)}{\delta_H(s' + t)e - \delta_L(s + t)b + c_L(s + t)} \equiv r^*$$

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

6.3 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 low type. Consider the first case. Then D's equilibrium utility for making the associated offer is $-p(\delta_L(s)b - c_L(s)) - (1 - p)\delta_H(s')e$. If D settles with identical types following the transfer, the expected probabilities remain the same. The equilibrium concession to the low type goes up, however, and the high type develops nuclear weapons faster. Thus, D's equilibrium utility equals $-p(\delta_L(s+t)b - c_L(s+t)) - (1 - p)\delta_H(s' + t)e$. 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 $-(\delta_H(s')b - c_H(s'))$. If it makes the transfer and settles with both types, its utility equals $-(\delta_H(s' + t)b - c_H(s' + t))$. Each constituent component is again smaller in the second case, so D cannot provide assistance that induces those bargaining strategies post-transfer either. $\square$

6.4 Proof of Lemma 4

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

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

6.5 Proof of Proposition 1

The following lemma will be useful:

Lemma 7. *There exists a critical value of q^* strictly less than .5 such that $r^* \in (\overline{r}, \underline{r})$ for all $q > q^*$, whereas $r^* \notin (\overline{r}, \underline{r})$ for all $q < q^*$.*

In words, if the strength of the signal is sufficiently high, then the posterior belief D infers following a low signal induces D to make the risky demand whereas the posterior belief D infers following a high signal induces D to make the safe demand. This follows immediately from three observations:

1. r^* is strictly bound between 0 and 1
2. $\underline{r}$ strictly increases in q and goes to 1 as q goes to .5
3. $\overline{r}$ strictly decreases in q and goes to 0 as q goes to .5

I now split the proof into two cases: $p < p^*$ and $p > p^*$. First, suppose $p < p^*$. Thus, without assistance, it proposes $\delta_H(s')b - c_H(s')$. If $q < q^*$, then Lemmas 4 and 7 show that D must not give assistance and propose that amount in equilibrium. If $q > q^*$, then D must check whether its payoff for providing assistance given the induced strategies exceeds its payoff for withholding assistance. D's payoff for the latter case is $-(\delta_H(s')b - c_H(s'))$.

The former case is more complicated. With probability $p(.5+q)$, R is the low type and D receives the signal indicating a low type. Because $q > q^*$, it proposes $\delta_L(s+t)b - c_L(s+t)$, and R accepts. D's payoff is the negative of that amount. With probability $(1-p)(.5-q)$, R is the high type and D receives the signal indicating a low type. Because $q > q^*$, D proposes an insufficient amount, causing R to reject and generating a payoff of $-\delta_H(s'+t)e$ for D. With probability $.5+q-2pq$, D receives the signal indicating a high type. Because $q > q^*$, D proposes $\delta_H(s'+t)b - c_H(s'+t)$, and R accepts regardless of its true type. D's payoff is the negative of that quantity. In all cases, D pays k.

In total, D prefers giving assistance if:

$$-p(.5+q)(\delta_L(s+t)b - c_L(s+t)) - (1-p)(.5-q)(\delta_H(s'+t)e)$$
$$-(.5+q-2pq)(\delta_H(s'+t)b - c_H(s'+t)) > -(\delta_H(s')b - c_H(s'))$$

Solving for q yields:

$$q > \frac{.5p(\delta_L(s+t)b - c_L(s+t)) + .5(1-p)(\delta_H(s'+t)e) + .5(\delta_H(s'+t)b - c_H(s'+t)) - (\delta_H(s')b - c_H(s')) + k}{(1-p)(\delta_H(s'+t)e) + (2p-1)(\delta_H(s'+t)b - c_H(s'+t)) - p(\delta_L(s+t)b - c_L(s+t))} \quad (1)$$

Thus, if $q > q^*$ and Line 1 holds, D gives assistance. Otherwise, it does not. Note that as k increases, the necessary threshold for q in Line 1 also increases. A large enough k implies that the cutpoint exceeds 1, and thus D would not give assistance even for a perfect signal.

Now suppose $p > p^*$. Then, in the absence of a transfer, D proposes $\delta_L(s)b - c_L(s)$. The low type accepts and the high type rejects. Thus, D's payoff equals $-p(\delta_L(s)b - c_L(s)) - (1-p)(\delta_H(s)e)$. Making the transfer generates the same payoff as the last case. As such, D prefers giving assistance if:

$$-p(.5+q)(\delta_L(s+t)b - c_L(s+t)) - (1-p)(.5-q)(\delta_H(s'+t)e)$$
$$-(.5+q-2pq)(\delta_H(s'+t)b - c_H(s'+t)) - k > -p(\delta_L(s)b - c_L(s)) - (1-p)(\delta_H(s)e)$$

Solving for q yields:

$$q > \frac{.5p(\delta_L(s+t)b - c_L(s+t)) + .5(1-p)(\delta_H(s'+t)e) + .5(\delta_H(s'+t)b - c_H(s'+t)) - p(\delta_L(s)b - c_L(s)) - (1-p)(\delta_H(s)e) + k}{(1-p)(\delta_H(s'+t)e) + (2p-1)(\delta_H(s'+t)b - c_H(s'+t)) - p(\delta_L(s+t)b - c_L(s+t))} \quad (2)$$

Thus, if $q > q^*$ and Line 2 holds, D gives assistance. Otherwise, it does not. Like before, note that as k increases, the necessary threshold for q in Line 1 also increases. A large enough k implies that the cutpoint exceeds 1, and therefore D would not give assistance even for a perfect signal.

6.6 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 low type accepts, and the high type rejects. The probability of proliferation is $1 - p$. With assistance, the probability of proliferation is the probability that R is high *and* the probability D receives a low signal given that R is high. 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.[32] For q below the threshold, the probability of proliferation equals $1 - p$ and is therefore constant 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. $\square$

6.7 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 the high type *and* the probability D receives a low signal given that R is high. 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

[32] "Sufficiently high" here can exceed .5, in which case the probability of proliferation remains constant at $1 - p$.

low and does if q is sufficiently high.[33] For q below the threshold, the probability of proliferation is 0 and therefore constant 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 (from 0 to a strictly positive amount), making the overall relationship nonmonotonic. $\square$

7 Extensions

To simplify notation, I assume that $\delta_H = \delta_L$ and $c_H = c_L$ throughout the extensions. Relaxing these does not change the results.

7.1 Monitoring Problems

In the main model, bargaining was *quid-pro-quo*—accepting x implies that R does not build. One may wonder whether negotiations can still work when this does not hold. I now show that they can as long as proliferation is sufficiently inefficient. Moreover, with uncertainty, the game still exhibits the same risk-return tradeoff that drives Propositions 1 and 2.

To formalize the monitoring problem, suppose that R's payoff for building now equals $(1-\delta(\bullet))x + \delta(\bullet)b - c(\bullet)$. That is, for the $1-\delta(\bullet)$ length of time, R can enjoy whatever benefits D has offered it before transitioning to the b benefits thereafter. This represents a worst-case scenario for D, as it does not observe R's decision (and thus retains its concession) until R's proliferation process is over. Accepting yields the same payoff as before.

First, I show that deals still work in this environment. Consider R's accept versus build decision. R prefers accepting if:

$$x \geq (1-\delta(\bullet))x + \delta(\bullet)b - c(\bullet)$$
$$x \geq b - \frac{c(\bullet)}{\delta(\bullet)}$$

To be willing to make this best possible offer, D must prefer it to allowing prolifer-

[33] Again, "sufficiently high" here can exceed .5, in which case the probability of proliferation remains constant at 0.

ation to transpire, or $x \leq \delta(\bullet)e$. Note this has not changed from before, as D does not suffer any additional negative consequences from R secretly building if D has not made any concessions at all. Thus, as long as $\delta(\bullet)e \geq b - \frac{c(\bullet)}{\delta(\bullet)}$ holds, agreements still work. D offers $b - \frac{c(\bullet)}{\delta(\bullet)}$ in the bargaining subgame, and R accepts.

The second question is whether the general logic of Propositions 1 and 2 apply to this setup. To see why it does, consider how D resolves its risk-return tradeoff. A type with a greater competency requires a larger offer to accept. As such, D can adopt one of two strategies. First, it can demand just enough to induce the low type to accept. This causes the high type to reject, holding onto the offered concession in the meantime. Second, it can demand just enough to induce the high type to accept. This also buys the low type's compliance.

Working out the expected payoffs for each choice, and using the non-transfer case as a starting point, D prefers the risky offer if:

$$-q\left(b - \frac{c(s)}{\delta(s)}\right) - (1-q)\left((1-\delta(s'))\left(b - \frac{c(s)}{\delta(s)}\right) + \delta(s)e\right) > -\left(b - \frac{c(s')}{\delta(s')}\right)$$

$$q > \frac{\delta(s')(e-b) + \frac{c(s)\delta(s')}{\delta(s)} - \left(\frac{c(s)}{\delta(s)} - \frac{c(s')}{\delta(s')}\right)}{\delta(s')(e-b) + \frac{c(s)\delta(s')}{\delta(s)}}$$

By analogous reasoning, obtaining the cutpoint on q for the transfer case simply requires adding t to the functions' arguments.

Thus, D's decision boils down to a standard risk-return tradeoff. As a result, the probability of proliferation can increase or decrease following the transfer, depending on its prior belief. For example, if q starts above the cutpoint, D would make the risky demand in the absence of a signal. But a transfer can generate a signal that causes it to make the safe demand instead, thereby decreasing the probability of proliferation. On the other hand, if q starts below the cutpoint, D would make the safe demand in the absence of a signal. But a transfer can generate a signal that causes it to make the risky demand instead, thereby increasing the probability of proliferation.

7.2 Sanctions

The main text noted that more proficient types find sanctions less threatening than less proficient types. In turn, more proficient types would want to signal that quality, which allows for the unraveling effect to play out if the recipient could accept or reject the assistance. I now formalize the logic.

Following theoretical and empirical work on sanctions and leader turnover (Marinov, 2005; Spaniel and Smith, 2015), suppose that attempting proliferation leads to sanctions that lower the probability of leader survival. Only a leader who survives to proliferation enjoys the benefits. To capture these incentives, suppose that for a given level of proficiency, it takes t periods to proliferate, δ is a traditional discount factor, and ϕ is the probability of survival in any given period; $1-\phi$ therefore captures the effectiveness of sanctions. Then, rewriting the $\delta(\bullet)$ function from the main model accordingly, that type's payoff for rejection is:

$$\phi^t \delta^t b - c(t)$$

Compare this to a less competent type that takes t' periods to proliferate. The more competent type has a greater proliferation payoff if:

$$\phi^t \delta^t b - c(t) > \phi^{t'} \delta^{t'} b - c(t')$$

$$c(t') - c(t) > b(\phi^{t'} \delta^{t'} - \phi^t \delta^t)$$

The left hand side is positive because less competent types have a greater cost than more competent types. Because $\phi^t \delta^t > \phi^{t'} \delta^{t'}$, the right hand side is negative. As such, with sanctions incorporated into the model, high proficiency types still have a greater proliferation payoff. This is because proficiency both allows nuclear weapons to be built sooner and risks fewer periods of leadership turnover in the meantime for any hazard rate consistent between the types.

As an additional note, one may wonder if this result would change if the opponent might choose differential sanctions (and thus different values of ϕ) depending on the competency of the recipient. In fact, provided that the types proliferate relatively fast—which the assumption that both types have positive proliferation values pushes them toward—this is true. An extreme case makes the intuition evident. Consider the incentive to sanction a country that can proliferate almost instantaneously. Then marshaling sanctions and causing long-term damage to one's economy serves little purpose. The sanctions barely have any time to take effect and unseat the leader. Sanctions instead have greater utility with slower types that internalize the risk of turnover to a non-trivial degree.

Furthermore, consider a game in which the opponent could adjust its sanctions capacity in every period once the recipient has decided to proliferate. Then a less proficient type must have a weakly lower probability of leader survival. This is because of two facts. First, failed sanctions in any previous period are sunk costs. And second, the sanctioning subgame for the more proficient type is a proper subgame of the sanctioning subgame for the less proficient type. In other words, consider the optimal sequence of sanction levels when a state is t periods away from nuclear weapons. The associated probability of survival is a more proficient type's probability of survival. A less proficient type must also survive through the optimal sequence of sanctions levels from periods t' to $t+1$ before proliferation. If no sanctions would take place over those periods, then the less proficient type has an equivalent probability of survival. Otherwise, the less proficient type is strictly less likely to survive.