Give Peace a (Second) Chance: A Theory of Nonproliferation Deals
Online Appendix

Muhammet A. Bas and Andrew J. Coe

Section 1 discusses the differences between our model and others of nuclear proliferation. Section 2 presents the propositions (and proofs thereof) needed to characterize what happens in the absence of a deal (i.e., the no-deal equilibrium). Section 3 presents proofs of the propositions stated in the main body of the paper, characterizing equilibrium deals. Section 4 presents the summary and formal analysis of the conditions under which early or late deals are viable. No-deal propositions are indexed with Roman numerals, while deal propositions have Arabic numbers, as in the main body of the paper. Section 5 presents additional details and sources for the empirical cases not discussed in the main body.

1 Differences from Other Models of Proliferation

The setup of our model differs importantly from other models of nuclear proliferation such as Debs and Monteiro (2014) (henceforth, DM). In DM, B first decides whether to invest, then a signal of his decision is sent, then A either attacks or not. Crucially, if B invests, there is no chance that his effort succeeds before A has a chance to attack. Hence, if the signal if perfectly informative, and A prefers attack to proliferation, then in equilibrium B refrains from investment in order to avoid certain attack, and no deal is possible. If the signal is not perfectly informative, then equilibrium features B randomizing over investment and A randomizing over attacking if the signal is that B has not invested.

Our model orders the moves differently and so produces different results. If B invests, A will subsequently have an opportunity to detect this and respond accordingly, *but not instantaneously.* In the time until A gets the signal, B might succeed with some (possibly very small) probability, regardless of how informative the signal is. As we will explain subsequently, regardless of the signal's informativeness, in the absence of a deal, B will invest in order to try to improve his bargaining power and A will wait until she believes his program is nearing success to resort to costly conflict. Unlike in DM, there is no randomization, and a deal is always needed if proliferation or costly conflict are to be avoided.

We view our setup as a step forward in the modeling of nuclear proliferation. DM assumes that preventive attack or proliferation (if there is no attack) occur immediately after investment, a useful simplification that yields important insights. However, proliferation is actually a dynamic phenomenon: even the fastest programs took years to succeed, and most programs are suspected and watched by other states for much of their lives. Our setup incorporates this fact. Realistically, the United States does not instantly detect the initiation of a nuclear weapons program. Even if it did, it cannot instantly resort to costly conflict to stop it. And even if it could, it would not do so in equilibrium because a program that has just begun poses no threat of imminent proliferation, so that the costs of conflict can and will be safely put off.

2 No-Deal Equilibrium

A "no-deal equilibrium" is defined to be a Perfect Bayesian Equilibrium in which A never makes a positive offer that is more generous than is necessary to avoid immediate costly conflict. More precisely, it is a PBE in which, at the beginning of every period, either B's continuation value is equal to his costly conflict value or A's offer in that period is $q = 1$. Intuitively, the only reason A would ever make such an offer is to induce B not to invest, in the context of a deal trading these concessions for non-investment.

We start with a simple lemma that determines equilibrium behavior once B has acquired nuclear weapons.

Lemma I. *Suppose that B has acquired nuclear weapons. In every period, A will offer $q = p_n + c_B$, and B will accept $q \leq p_n + c_B$. No costly conflict will occur.*[1]

Proof. Subgame perfection implies that B will accept any $q < p_n + c_B$, because rejecting it yields his costly conflict value, while accepting it and initiating costly conflict in the next round guarantees a higher payoff. Because of this, A would strictly prefer offering any $q \in (p_n - c_A, p_n + c_B)$ to costly conflict. For any such q that is less than $p_n + c_B$, there is a higher q that both A and B would strictly prefer to costly conflict. Thus, in equilibrium A makes the offer that renders B indifferent to costly conflict ($q = p_n + c_B$) and B accepts this or any higher offer and rejects any lower offer. □

Proposition I. *Suppose that B is in the second stage and this is common knowledge. In any no-deal equilibrium of this subgame, B always invests, given the chance. If $p - c_A + \frac{\delta\lambda}{1-\delta}(p - p_n) < 1 + \frac{\delta\lambda}{1-\delta}(c_A + c_B)$, there is peace and eventual proliferation. Otherwise, there is immediate costly conflict and no proliferation.*

Proof. Suppose that, in some period prior to B obtaining nuclear weapons, his equilibrium continuation value is at least $V_n^B \equiv (1-p_n-c_B)/(1-\delta)$. This implies that A's own value is at most $1/(1-\delta) - V_n^B = (p_n+c_B)/(1-\delta)$. But then A could profitably deviate by offering some $q \in (p_n + c_B, \max\{p + c_B, 1\})$ in all rounds prior to proliferation, which subgame perfection requires B to accept. This violates the supposition of equilibrium, so B's pre-proliferation continuation value must always be less than V_n^B.

[1] We assume throughout that $p_n + c_B < 1$. If this did not hold, then B would have no bargaining power, with or without nuclear weapons, and so no incentive to acquire them, and A would have no reason to try to stop B from getting them. We discard this uninteresting case.

Since, in a no-deal equilibrium, A does not react to signals of investment, an investment that fails gives B the same continuation value in the next period as not investing, and this value is less than V_n^B. But the investment succeeds with positive probability, yielding a next-period value of V_n^B, and so B always strictly prefers to invest, given the chance.

In any peaceful no-deal equilibrium, B must receive at least $\underline{V}^B \equiv \max\left\{\frac{1-p-c_B}{1-\delta}, \frac{\delta\lambda V_n^B}{1-\delta(1-\lambda)}\right\}$, since if he received less than the first term, he would do better by initiating costly conflict, and he cannot receive less than the second term even if A offers $q = 1$ in every round before B obtains nuclear weapons. This means that A's continuation value is at most $1/(1-\delta) - \underline{V}^B$. If this is less than A's costly conflict value $W^A \equiv (p - c_A)/(1 - \delta)$, then the unique no-deal equilibrium features immediate costly conflict. It is easily seen that this can only occur when $\underline{V}^B = \frac{\delta\lambda V_n^B}{1-\delta(1-\lambda)}$. Re-arranging the inequality $1/(1-\delta) - \frac{\delta\lambda V_n^B}{1-\delta(1-\lambda)} < W^A$ leads to the condition for costly conflict given in the statement of the proposition.

Suppose instead that $1/(1-\delta) - \underline{V}^B > W^A$ and that costly conflict occurs in equilibrium. In the period in which it occurs, subgame perfection requires B to accept any

$$q \in \left(p - c_A + \frac{\delta\lambda}{1-\delta}(p - p_n - c_A - c_B), \min\left\{1, p + c_B + \frac{\delta\lambda}{1-\delta}(p - p_n)\right\}\right),$$

since taking this and initiating costly conflict in the next round yields a better value for B than doing so now, given that B is investing. The supposed condition implies that this range is non-empty, and any offer in it would also be better for A than costly conflict, so at least one player has a profitable deviation to not initiating costly conflict, and thus it cannot occur in equilibrium. Since, for any offer in the above range, A could make a slightly less generous offer that B would still accept, in equilibrium A offers $q = \min\left\{1, p + c_B + \frac{\delta\lambda}{1-\delta}(p - p_n)\right\}$, and B accepts this or any higher offer and rejects any lower offer. □

Proposition II. *Suppose that the second-stage-known equilibrium is immediate costly conflict. In any no-deal equilibrium of the game, B always invests, given the chance.*

Proof. We first establish equilibrium behavior in the subgames in which B's program has reached the second stage. We then turn to the first-stage subgames. For convenience, we use "$B1$" and "$B2$" to refer to B whenever his program is in the first or second stage, respectively.

We start by establishing some properties of any no-deal equilibrium that will be used in later arguments. Observe that B's continuation value from the beginning of any period must be at least his costly conflict value W^B. If it is not, B could profitably deviate by rejecting A's offer and thereby causing costly conflict. Consequently, if the current round is peaceful and B does not have nuclear weapons, B can invest and thereby guarantee himself an expected next-round continuation value greater than his costly conflict value. If the investment fails, he will receive at least his costly conflict value in the next round, but it will succeed (in the sense of B getting nuclear weapons) with positive probability and, by Lemma I, yield a value in the next round of $V_n^B > W^B$. It follows that, in any round prior to B acquiring nuclear weapons, the minimum offer B would accept gives him less than his per-period costly conflict value. Thus, in a no-deal equilibrium, A's offer will always give B less than $1 - p - c_B$, if it is positive, and zero otherwise.

As supposed in the statement of the proposition, if A knows that B's program has reached the second stage, then A immediately attacks or abandons, and if he did not then B would invest. So consider any prior period in which A faces $B2$, but has yet to detect this, and suppose that $B2$'s continuation value is at least V_n^B. For this to be true, it must be that A offers a $q \leq p_n + c_B$ in that or some later period before proliferation. But this violates our observation that A's offer is always at least $\min\{p + c_B, 1\}$, so $B2$'s continuation value must always be less than V_n^B. Then the same argument as used in the proof of Proposition I applies, and $B2$ will always invest, given the chance.

Moreover, the supposition that a second-stage signal leads to costly conflict implies that $B2$ strictly prefers investment to costly conflict, even if A offers him nothing. To see this, note that failed investment must yield a next-period value of at least W^B, so that the present value of investment is at least $\delta\left[\lambda V_n^B + (1-\lambda)W^B\right]$, and Proposition I implies this is greater than W^B. Thus, once B has reached the second stage, he will always accept any offer from A and invest. It follows that, if A ever has to give B more than nothing, it is because this is required to satisfy $B1$. This in turn implies that, in equilibrium, A will never make an offer that $B1$ would reject. To see why, note that rejection of the offer would bring costly conflict, while acceptance would enable A to infer that B's program had reached the second stage. If there is any possibility of the latter, A would do better by attacking/abandoning rather than making the offer; otherwise, the offer makes costly conflict certain and so we treat it as equivalent to attacking/abandoning.

Now consider any period (call it t) in which B's program is in the first stage, and suppose by way of contradiction that there is a no-deal equilibrium in which $B1$ would not invest at t, given the chance to do so. $B1$'s continuation values of investing (I) and not (NI) and $B2$'s continuation value are:

$$\begin{aligned}
V_t^{B1}(\text{I}) &= 1 - q_t + \delta\epsilon\left[\lambda V_n^B + (1-\lambda)\left(\sigma W^B + (1-\sigma)V_{t+1}^{B2}\right)\right] + \delta(1-\epsilon)\left[\sigma V_0^B + (1-\sigma)V_{t+1}^{B1}\right] \\
V_t^{B1}(\text{NI}) &= 1 - q_t + \delta\left[\sigma V_0^B + (1-\sigma)V_{t+1}^{B1}\right] \\
V_t^{B2} &= 1 - q_t + \delta\left[\lambda V_n^B + (1-\lambda)\left(\sigma W^B + (1-\sigma)V_{t+1}^{B2}\right)\right]
\end{aligned}$$

where V_{t+1}^{B2} is $B2$'s continuation value from the next period, given that he did not get nuclear weapons this period, V_0^B is $B1$'s continuation value from the next period, given that A received a signal that his program remained at the first stage, and V_{t+1}^{B1} is $B1$'s value from the next period given that A received no stage signal.

Since $B1$ does not invest at t, it must be that $V_t^{B1}(\text{I}) \leq V_t^{B1}(\text{NI})$, which is equivalent to

$V_t^{B2} \leq V_t^{B1}(\cdot)$, or:

$$\lambda V_n^B + (1-\lambda)\left(\sigma W^B + (1-\sigma)V_{t+1}^{B2}\right) \leq \sigma V_0^B + (1-\sigma)V_{t+1}^{B1}. \tag{1}$$

By expansion of V_{t+1}^{B2}, the left hand side of 1 is at least:

$$\underline{V_L} \equiv V_n^B \sum_{i=0}^{\infty} \left[\Pi_{j=1}^{i}(1-p_j)\right] \delta^i \mu^i \lambda + W^B \sum_{i=0}^{\infty} \left[\Pi_{j=1}^{i-1}(1-p_j)\right] \delta^i \left[p_i + (1-p_i)\mu^i(1-\lambda)\sigma\right]$$

where $\mu \equiv (1-\lambda)(1-\sigma)$ and p_j is the probability that A attacks/abandons at the beginning of period $t+j$, given that B does not have nuclear weapons and A has not received a stage signal since at least period t. The "at least" follows from the fact that this expansion neglects any value B would attain from A's offers. Henceforth, we will abbreviate $\Pi_{j=1}^{i}(1-p_j)$ as P_i.

On the right hand side of 1, V_0^B must be less than:

$$\overline{V_0^B} \equiv V_n^B \sum_{i=0}^{\infty}(1-\epsilon)^i \epsilon \sum_{j=0}^{\infty} \mu^j \lambda + W^B \sum_{i=0}^{\infty}(1-\epsilon)^i \epsilon \sum_{j=0}^{\infty} \mu^j (1-\lambda)\sigma = V_n^B \psi + W^B(1-\psi)$$

where $\psi = \frac{\lambda}{\lambda+\sigma-\lambda\sigma}$. ψ is the probability that the game will end with B's proliferation, assuming that A would never attack/abandon B unless he received a signal that B's program had reached the second stage, and is thus an upper bound on the equilibrium probability of proliferation. Since the per-period payoffs associated with V_n^B exceed those associated with costly conflict, which exceed those associated with A's pre-proliferation offers, and the formula ignores discounting due to any delay in costly conflict or proliferation occurring, this formula must be greater than V_0^B.

We deal with the second value on the right hand side of 1 (V_{t+1}^{B1}) in two cases. First suppose that, given the chance, $B1$ will not invest again after t, unless A receives a signal of his stage. For this to be in equilibrium, it must be that $V_{t'}^{B1} > W^B$ for every $t' \geq t$ in which A has not yet received a stage signal. (If instead $V_{t'}^{B1} \leq W^B$, there would be a profitable deviation to investment, since investing is strictly preferred to not investing whenever $V_{t'}^{B2} > V_{t'}^{B1}$, and we observed earlier that it must always be true that $V_{t'}^{B2} > W^B$.) This in turn implies that the offer $q_{t'}$ associated with period t' must be 1, so that A offers nothing to B. (If instead $q_{t'} < 1$, then this cannot be a no-deal equilibrium, since A could make a less generous offer and still avoid costly conflict in that period with certainty.) Then by expansion of V_{t+1}^{B1}, and substituting $\overline{V_0^B}$ for V_0^B, the right hand side of 1 is less than:

$$\overline{V_R} \equiv \overline{V_0^B} \sum_{i=0}^{\infty} P_i \delta^i (1-\sigma)^i \sigma + W^B \sum_{i=0}^{\infty} P_{i-1} \delta^i p_i$$

Using our lower and upper bounds on the two sides, 1 thus implies that $\underline{V_L} < \overline{V_R}$. Subtracting $W^B \sum_{i=0}^{\infty} P_{i-1}\delta^i p_i$ from both sides and collecting terms, we have:

$$\sum_{i=0}^{\infty} P_i \delta^i \mu^i \left[\lambda V_n^B + (1-\lambda)\sigma W^B\right] < \sum_{i=0}^{\infty} P_i \delta^i (1-\sigma)^i \sigma \left[\psi V_n^B + (1-\psi) W^B\right]$$
$$\Leftrightarrow (\lambda + \sigma - \lambda\sigma) \sum_{i=0}^{\infty} P_i \delta^i \mu^i < \sigma \sum_{i=0}^{\infty} P_i \delta^i (1-\sigma)^i$$

This inequality is false. To see why, temporarily set aside the factors of $P_i\delta^i$ from the two series. Since $(\lambda+\sigma-\lambda\sigma)\sum_{i=0}^{\infty}\mu^i = \sigma\sum_{i=0}^{\infty}(1-\sigma)^i = 1$, and $\mu^i \leq (1-\sigma)^i$ for all i, it must be that each partial sum of the simplified left hand side exceeds the corresponding partial sum of the simplified right hand side. Returning to the unsimplified series, because $P_i\delta^i$ is at most one and never increases in i, it must shrink later terms in each series at least as much as earlier terms. Thus, both unsimplified series will converge, but the one on the right cannot possibly exceed the one on the left, since the latter converges more quickly relative to its simplified version and is therefore reduced less by the presence of $P_i\delta^i$.

This contradiction eliminates the possibility that in a no-deal equilibrium $B1$ would not invest from some period onward until A received a stage signal. So now suppose that $B1$ does not invest at period t, and, in the absence of a stage signal, waits until the period $t' > t$

to invest again. It is sufficient to show that 1 is contradicted when $t' = t + 1$, because for any larger t', $B1$ will not invest at period $t'' = t' - 1$, giving rise to the same contradiction except at t'' instead of t. Since $B1$ invests at t', it must be that $V_{t'}^{B1} \leq V_{t'}^{B2}$, so 1 implies:

$$\lambda V_n^B + (1-\lambda)\sigma W^B + \mu V_{t'}^{B2} < \sigma \overline{V_0^B} + (1-\sigma)V_{t'}^{B2}$$
$$\Leftrightarrow \psi V_n^B + (1-\psi)W^B = \overline{V_0^B} < V_{t'}^{B2}$$

This inequality is false. To see why, note that the argument establishing $\overline{V_0^B}$ as an upper bound for V_0^B implies that it is also an upper bound for B's continuation value in *any* period prior to proliferation, regardless of stage. Thus, $V_{t'}^{B2} \leq \overline{V_0^B}$.

Thus, regardless of $B1$'s subsequent investment behavior, it cannot be in (a no-deal) equilibrium for him to not invest at period t. Since t is arbitrary, the result is established. □

Proposition III. *In any no-deal equilibrium, each consecutive null signal of B's stage increases A's estimate of the probability that B's program has reached the second stage. If* $\epsilon \geq \lambda$, *this estimate will converge to 1; otherwise it will converge to* $\frac{\epsilon - \epsilon\lambda}{\lambda - \epsilon\lambda} < 1$.

Proof. Since B always invests, signals of his investment are irrelevant to A's estimate of his stage. After i consecutive null stage signals since A was last certain that B was in the first stage, the probability that B remains in the first stage is just $(1-\epsilon)^i$. The probability that B has reached the second stage is the sum of the probabilities that he reached the second stage at any given point since A was last certain of his stage, and then did not subsequently acquire nuclear weapons, or $\sum_{j=1}^{i}(1-\epsilon)^{i-j}\epsilon(1-\lambda)^j$. Since B has not obtained nuclear weapons (recall we assumed this would immediately become common knowledge), these are the only two possibilities, and A's estimate is:

$$\frac{\sum_{j=1}^{i}(1-\epsilon)^{i-j}\epsilon(1-\lambda)^j}{(1-\epsilon)^i + \sum_{j=1}^{i}(1-\epsilon)^{i-j}\epsilon(1-\lambda)^j}$$

To see how the estimate converges, factor $(1-\epsilon)^i$ out of numerator and denominator and cancel these (since $\epsilon < 1$) to obtain $\frac{\epsilon \sum_{j=1}^{i} \alpha^j}{1+\epsilon \sum_{j=1}^{i} \alpha^j}$, where $\alpha = (1-\lambda)/(1-\epsilon)$. If $\epsilon \geq \lambda$, then $\alpha \geq 1$, and the estimate clearly converges to 1 as $i \to \infty$. Otherwise, $\alpha < 1$, and each sum converges to $\alpha/(1-\alpha) = (1-\lambda)/(\lambda-\epsilon)$, so that the estimate converges to $(\epsilon - \epsilon\lambda)/(\lambda - \epsilon\lambda)$. □

Proposition IV. *Suppose that the second-stage-known equilibrium is immediate costly conflict. In any no-deal equilibrium of the game, A tolerates B's program until she becomes sufficiently confident that it has reached the second stage, and then attacks/abandons.*

Proof. Proposition II establishes B's optimal behavior in a no-deal equilibrium, and Proposition III establishes A's beliefs in such an equilibrium. Here, we show that A's best response to B's strategy must be of the form given in the statement of the proposition.

First observe that, in equilibrium, any offer A makes must render $B1$ indifferent between acceptance and rejection, or, if that is infeasible, must give B nothing (i.e., $q = 1$). It was shown in the proof of Proposition II that it is always possible for A to satisfy B: $B1$ would strictly prefer to accept $q \leq \min\{p + c_B, 1\}$, and $B2$ would strictly prefer to accept any offer whatsoever. Since $B1$'s continuation value of acceptance varies continuously in A's current offer, either there is an offer that renders $B1$ indifferent between costly conflict and peace, or B will accept any offer regardless of stage. If there is an offer that renders $B1$ indifferent between costly conflict and peace, it cannot be in equilibrium for A to offer more: by Proposition II, such an offer has no effect on B's behavior, so that A's generosity is wasted. It was also shown in the proof of Proposition II that it cannot be in equilibrium for A to offer less (A would strictly prefer to initiate costly conflict). Thus, equilibrium requires that B accept an offer that renders $B1$ indifferent between acceptance and rejection, and that A's offer must be this one, when it is feasible. Similarly, if this offer is not feasible, then A's offer must be $q = 1$, and B must accept it. These requirements pin down the offers A will make in equilibrium.

Starting from any subgame prior to costly conflict or proliferation, and up to the occurrence of a first-stage signal that would reset A's estimate of the probability he faces $B2$ to 0, A's strategy consists of a vector of offers $\vec{q}$ to be made after each subsequent consecutive null signal of B's stage, and a vector of probabilities that A will attack/abandon after each consecutive null signal, $\vec{\pi}$. After receiving i consecutive null stage signals, let V_i^A be A's continuation value of making an offer just sufficient to satisfy $B1$ in that round (or $q = 1$ if this is infeasible) rather than initiating costly conflict.

We will show that, in any subgame of any no-deal equilibrium, V_i^A must strictly decrease in i. This implies that there may come a point at which A has received enough consecutive null signals that his estimate of the probability he faces $B2$ is high enough to merit costly conflict rather than tolerating further risk of proliferation. It also implies that once A has reached this threshold, he will attack/abandon with certainty after any higher number of consecutive null signals, in accordance with the proposition.

We restrict consideration to equilibria in which there exists some $\bar{\imath}$ such that, for all $i \geq \bar{\imath}$, $\pi_i \in \{0, 1\}$. That is, we require that in equilibrium, once A has received sufficiently many consecutive null signals of B's stage, then A will not randomize over whether to initiate costly conflict in this period, or after any number of additional consecutive null signals. This restriction simplifies the proof, but it also rules out some empirically implausible equilibria. By Proposition III, as additional consecutive null signals are received, A's estimate of the probability that he faces $B2$ converges, so that any strategy excluded by this restriction would require A to randomize over costly conflict at some point arbitrarily close to his estimate's limit. But it can be shown that, if there is an equilibrium of the subgame starting from that point, in which A would randomize in that period, then there is a Pareto-superior equilibrium in which A would not attack/abandon in that period or after any further consecutive null signals. (Not attacking/abandoning in that period raises the total value of the subgame, rendering B willing to agree to a less generous offer and leading A to strictly prefer making this offer to costly conflict.)

With this restriction in place, we start from a subgame occurring after A has received $i > \bar{\imath}$ consecutive null signals, to show that $V_{i-1}^A > V_i^A$ in any no-deal equilibrium. We will make use of the following general form of the players' continuation values when A makes an offer rather than initiating costly conflict after i consecutive null signals.

$$
\begin{aligned}
V_i^{B2} &= 1 - q_i + \delta \left[\lambda V_n^B + (1-\lambda)\left[\sigma W^B + (1-\sigma) V_{i+1}^{B2}\right]\right] \\
V_i^{B1} &= 1 - q_i + \delta\epsilon \left[\lambda V_n^B + (1-\lambda)\left[\sigma W^B + (1-\sigma) V_{i+1}^{B2}\right]\right] + \delta(1-\epsilon)\left[\sigma V_0^B + (1-\sigma) V_{i+1}^{B1}\right] \\
V_i^A &= q_i + \delta\left[\rho_i \epsilon + (1-\rho_i)\right]\lambda V_n^A + \delta\left[\rho_i\epsilon + (1-\rho_i)\right](1-\lambda)\sigma W^A \\
&\quad + \delta\rho_i(1-\epsilon)\sigma V_0^A + \delta\left[\rho_i(1-\epsilon) + \rho_i\epsilon(1-\lambda) + (1-\rho_i)(1-\lambda)\right](1-\sigma)\tilde{V}_{i+1}^A
\end{aligned}
$$

where $V_n^A \equiv \frac{p_n + c_B}{1-\delta}$ is A's continuation value once B has acquired nuclear weapons, V_0^A is A's continuation value once he receives a signal that B's program remains in the first stage, $\tilde{V}_{i+1}^A$ is A's continuation value in equilibrium after receiving $i+1$ null signals, and ρ_i is A's estimate of the probability that B's program remains in the first stage after receiving i consecutive null signals since the last first-stage signal (or the start of the game).

Consider the probabilities that A will initiate costly conflict after i and $i+1$ consecutive null signals: (π_i, π_{i+1}). We divide the possible values of this ordered pair into five cases to be analyzed in turn:

1. $(\pi, 1)$: The only possible differences in the equations for B's continuation values at $i-1$ and at i are in the offers A will make and the values B will receive in the subsequent round, in the absence of a non-null signal or successful acquisition of nuclear weapons. Since, in the absence of those events, B will receive his costly conflict value after $i+1$ consecutive null signals with certainty, but will receive at least this value after i consecutive null signals (and more if he is at or reaches the second stage), it follows that the discounted terms of V_{i-1}^{B1} will be at least as large as those of V_i^{B1}, and hence

that the offer A must make to satisfy $B1$ after $i-1$ signals will be no more generous than that required after i signals, so that $q_{i-1} \geq q_i$. Since $\tilde{V}_i^A = \tilde{V}_{i+1}^A = W^A$ (using the fact that randomizing over initiating costly conflict is in equilibrium only if both attacking/abandoning and not yield the same continuation value), the equations for V_{i-1}^A and V_i^A differ only in their estimates of the probability that B is in the first stage and in the offer A must make. By Proposition III, $\rho_{i-1} > \rho_i$. So, relative to V_{i-1}^A, V_i^A has increased transition probabilities to V_n^A and W^A and decreased transition probabilities to V_0^A and $\tilde{V}_{i+1}^A$. Since $V_n^A < W^A \leq \left\{V_0^A, \tilde{V}_{i+1}^A\right\}$, and $q_{i-1} \geq q_i$, then it must be that $V_{i-1}^A > V_i^A$.

2. $(0, \pi)$, with $0 \leq \pi < 1$: Let $j \geq i+1$ be the first number of consecutive null signals larger than i at which $\pi_j > 0$, if any such number exists. (We deal with the case where it does not below.) If it does, then by reasoning similar to that for the previous case, the discounted terms of V_{j-2}^{B1} will be at least as large as those of V_{j-1}^{B1}: only the latter transitions to a positive probability of costly conflict in the absence of a non-null stage signal or proliferation, and since every additional consecutive null signal brings a probability of costly conflict at least $0 = \pi_{j-1}$, the overall probability that the game will end in costly conflict is higher starting from $j-1$ than from $j-2$. Thus, it must be that $q_{j-2} \geq q_{j-1}$. The same argument from in the previous case for comparing V_{j-2}^A to V_{j-1}^A applies, with the exception that these two values also differ in the value of the continuation game A faces in the absence of a non-null stage signal or proliferation. At $j-2$, the absence of these events will lead to a value of V_{j-1}^A, which must be at least W^A in equilibrium, whereas at $j-1$, it leads to $V_{j-2}^A = W^A$, so that it must be that $V_{j-2}^A > V_{j-1}^A$. Then, by induction, it must be that the discounted terms of V_{i-1}^{B1} are at least as large as those of V_i^{B1}, so that $q_{i-1} \geq q_i$, $V_i^A > V_{i+1}^A$, and thus $V_{i-1}^A > V_i^A$.

3. $(\pi, 0)$, with $\pi > 0$: This pair cannot occur in equilibrium. The previous case implies

that, whatever the value of π_{i+2}, since $\pi_{i+1} = 0$, it must be that $V_{i+1}^A \geq W^A$, and since $V_{i+1}^A < V_i^A$, it must be that $V_i^A > W^A$, implying that $\pi > 0$ cannot be in equilibrium.

4. $\pi_i > 0$, $0 < \pi_{i+1} < 1$: This pair cannot occur in equilibrium. The previous case implies that $\pi_{i+2} \neq 0$; if it is equal to 1, the first case implies that $\pi_i > 0$ cannot be in equilibrium. Finally, if $\pi_{i+2} \in (0, 1)$, then let $j > i + 2$ be the first larger number of consecutive null signals for which π_j is either 0 or 1; j exists by virtue of our restriction on equilibria strategies, and the previous case implies that $\pi_j = 1$. But then the first case above implies that $\pi_{j-2} = 0$, and the second case implies that $\pi_{j-3} = 0$ if it exists, and so on all the way back to π_{i+1}, so that this last value cannot be positive in equilibrium.

5. The excluded possibility in the second case above is that $\vec{\pi}$ may be of the form $(\pi, 0, 0, \dots)$, so that if A did not attack/abandon at the first opportunity, then A would not do so no matter how many additional consecutive null signals he received. Because $B1$'s continuation value varies continuously in each component of $\vec{q}$, and every future round until a reset due to a first-stage signal looks the same to $B1$, except for possible variation in future offers from A, then there must be a constant offer q^* that is just sufficient to satisfy $B1$ in every round, or else $\vec{q} = q^* = 1$ is sufficient.

   Observe that any vector other than $\vec{q} = q^*$ cannot be in equilibrium. Obviously this is true for any other constant vector, which must either be too generous or too stingy to $B1$. So consider any non-constant vector. If every component of this vector is at least q^*, then $B1$ will not accept any of the offers that are above q^*, since the value $B1$ receives from this vector starting from that point must be less than the value he receives from the constant vector of q^*, and hence less than his costly conflict value. Similarly, if every component is at most q^*, then every offer less than q^* is too generous, since at each such point B will receive strictly greater than his costly conflict value. In either case, A could profitably deviate by changing his offer to q^* or initiating costly conflict. Thus, any non-constant vector that is in equilibrium must have components above, and components below, q^*. If, say, the ith component q_i is greater than q^*, then at least one component subsequent to q_i must be less than q^* to "make up the difference" to $B1$ and keep his continuation value at i at least equal to his costly conflict value. Because $B1$ discounts subsequent offers, the sequence of differences between q^* and subsequent offers below q^* must have a discounted present value to $B1$ of at least $(q_i - q^*)/\delta$. Because these decreases must be made up by subsequent offers less generous than q^*, the sequence of increases above q^* in these latter subsequent offers must have a discounted present value of at least $(q_i - q^*)/\delta^2$. By repeating the argument, the discounted present value of needed changes in subsequent offers from q^* can be made arbitrarily large, but of course the whole value of the game is finite—at most $1/(1-\delta)$—so a non-constant vector cannot be in equilibrium.

This implies that A's continuation value, given that B's program remains in the first stage, does not differ across rounds, since A will make the same offer in every round, $B1$ will accept it and invest, and the transition probabilities to the second stage, proliferation, costly conflict, or a reset due to a first-stage signal are the same; call this value V_1^A. Similarly, A's continuation value, given that B's program is in the second stage, does not differ across rounds and is denoted V_2^A. Thus, we have $V_i^A = q^* + \rho_i V_1^A + (1-\rho_i)V_2^A$ for all i. We know that $V_2^A < W^A$, by the presumption that A would attack/abandon if he knew he faced $B2$. Since $\vec{\pi} = 0$, equilibrium requires that $V_i^A \geq W^A$; this in turn implies that $V_1^A > V_2^A$. By Proposition III, ρ_i is strictly decreasing in i, so it follows that V_i^A is strictly decreasing in i.

This is enough to establish the result. Starting from any component of $\vec{\pi}$, suppose it is 0. Then the cases above imply that the previous component, and every one preceding it, must also be 0, and that V_i^A strictly increases as we move back to fewer consecutive null signals.

Suppose the starting component is instead some $\pi \in (0,1)$. Then the previous component, and every one preceding it, must be zero, and V_i^A again strictly increases in decreasing i. Finally, suppose the starting component is 1. Then the preceding component is either 1, $\pi \in (0,1)$, or 0; for all three possibilities, V_i^A strictly declines as we move back. We can then move the starting component back one consecutive null signal, and repeat. Thus, V_i^A must be strictly decreasing in i, and $\vec{\pi}$ must be non-decreasing in i, with at most one component that is strictly between 0 and 1, in any no-deal equilibrium. □

3 Proofs of Propositions on Deal Equilibria

Proposition 1

In any deal equilibrium of a given subgame, B must prefer not investing to investing at the first opportunity to do so. Let the equilibrium continuation values from this opportunity be V_d^A and V_d^B for the two players, and the continuation values starting from this subgame be V_{nd}^A and V_{nd}^B if B deviates to investing, this is detected by A, and the players thereafter act as prescribed by the worst possible no-deal equilibrium for each player.[2] Let $V_d \equiv V_d^A + V_d^B$ and $V_{nd} \equiv V_{nd}^A + V_{nd}^B$.

We will show that if a deal equilibrium exists, it must be that $V_d > V_{nd}$. Since the occurrence of costly conflict is the only way that value can be destroyed in the game, this inequality means that a deal exists only if the no-deal equilibrium has a positive probability of costly conflict.

B's expected payoff from deviating is a convex combination of V_n^B (if his investment suc-

[2] For the subgame in which B's program is known to be in the second stage, Proposition I implies that V_{nd}^A and V_{nd}^B are unique. For the earlier subgames, in which A is not certain that B's program has reached the second stage, multiple no-deal equilibria may exist. These equilibria differ only in how long A waits before becoming suspicious enough that B's program has reached the second stage to act. The worst of these for both players is the one in which A acts soonest. This is the worst equilibrium for B because it gives his program the fewest chances to succeed before A acts. It is the worst for A, because otherwise in a (worse for A) longer-wait equilibrium, A would have a profitable deviation to attacking sooner.

ceeds and he obtains nuclear weapons), V_{nd}^B, and V_d^B, the last being realized if B's deviation is undetected and B reverts back to not investing. Thus, for some $\alpha, \beta \in [0,1]$ and $\gamma \equiv \frac{\alpha}{\alpha+\beta}$, we can write $V_d^B = \alpha V_n^B + \beta V_{nd}^B + (1-\alpha-\beta)V_d^B = \gamma V_n^B + (1-\gamma)V_{nd}^B$. Since we must have $V_{nd}^B < V_n^B$, it follows that $V_d^B > V_{nd}^B$.

Similarly, the deal should give A at least its value from the no-deal equilibrium. Together, $V_d^B > V_{nd}^B$ and $V_d^A \geq V_{nd}^A$ imply that $V_d > V_{nd}$, proving the result. □

Proposition 2

Let $V_{d,ns}^A$ and $V_{d,ns}^B$ be the continuation values for A and B of an early deal, made when A believes B's program is nearing success (denoted the "NS" deal), and $V_{d,u}^A$ and $V_{d,u}^B$ be their values for a middle deal, made when A does not yet believe that B's program is nearing success (denoted the "U" deal). We will show that the equilibrium constraints on these values are tighter for the middle deal. This implies that the conditions will be least restrictive either for an early or a late deal.

Let $B2$ refer to player B when his program is currently in the second stage of development, and $B1$ analogously for the first stage. In order for a deal to be in equilibrium, each player (and either type of B, corresponding with program stage) must prefer the deal to reneging. Observe that we need only check B's deviation to investing in a nuclear weapons program. B's only other possible deviation is to rejecting the offer from A under the deal, which would bring costly conflict. The constraint that the deal must be better than costly conflict for B is the same, no matter what A believed about B's stage under the deal, and hence is irrelevant to the comparison. Below, we also use the fact that Assumption 1 guarantees that if a deal fails when A believes B's program to be nearing success, A will immediately attack/abandon, as this is the worst possible equilibrium punishment for B's cheating.

For $B2$, we have the constraints:

$$\begin{aligned} V_{d,ns}^B &\geq \lambda V_n^B + (1-\lambda)\left[\tau\sigma W^B + \tau(1-\sigma)W^B + (1-\tau)V_{d,ns}^B\right] \\ V_{d,u}^B &\geq \lambda V_n^B + (1-\lambda)\left[\tau\sigma W^B + \tau(1-\sigma)V_{nd,u}^{B2} + (1-\tau)V_{d,u}^B\right], \end{aligned}$$

where $V_{nd,u}^{B2}$ is the continuation value $B2$ will get if he is caught pursuing nuclear weapons but his stage remains unknown to A, and A does not believe the program to be nearing success. Observe that this possibility is the only difference between the two inequalities. But we must have $V_{nd,u}^{B2} \geq W^B$, because after getting caught, B could reject any offer A made and receive W^B. Thus, $B2$'s constraint under the NS deal is no tighter than under the U deal.

For $B1$, the constraints are:

$$\begin{aligned} V_{d,ns}^B \geq \quad & \epsilon\lambda V_n^B + \epsilon(1-\lambda)\left[\tau\sigma W^B + \tau(1-\sigma)W^B + (1-\tau)V_{d,ns}^B\right] \\ & +(1-\epsilon)\left[\tau\sigma V_{nd,c}^{B1} + \tau(1-\sigma)W^B + (1-\tau)V_{d,ns}^B\right] \\ V_{d,u}^B \geq \quad & \epsilon\lambda V_n^B + \epsilon(1-\lambda)\left[\tau\sigma W^B + \tau(1-\sigma)V_{nd,u}^{B2} + (1-\tau)V_{d,u}^B\right] \\ & +(1-\epsilon)\left[\tau\sigma V_{nd,c}^{B1} + \tau(1-\sigma)V_{nd,u}^{B1} + (1-\tau)V_{d,u}^B\right], \end{aligned}$$

where $V_{nd,u}^{B1}$ is defined analogously to that value for $B2$, and $V_{nd,c}^{B1}$ is the continuation value $B1$ will receive when A is certain B's program is first-stage. As with $B2$, the only difference in the two constraints is when B is caught cheating, but his stage remains unknown to A. By the same argument as above, we must have $V_{nd,u}^{B1} \geq W^B$, so $B1$'s constraint under the NS deal is no tighter than under the U deal.

Finally, for A, the constraints are:

$$V_{d,ns}^A \geq W^A \qquad\qquad V_{d,u}^A \geq V_{nd,u}^A,$$

where $V^A_{nd,u}$ is the continuation value A will receive upon reneging on the deal by making a less-generous offer to B than the deal calls for. Just as for B, the constraint that A must prefer the deal to costly conflict is the same for either type of deal, and just as above, this implies that $V^A_{nd,u} \geq W^A$. Thus, A's constraint under the NS deal is no tighter than under the U deal.

It is shown in the proof of Proposition IV that A's equilibrium continuation value in the absence of a deal strictly declines as he grows more confident that B's program has reached the second stage. This implies that $V^A_{nd,u}$ is strictly greater than W^A, which establishes that the constraints for a U deal must be strictly tighter than those for an NS deal. □

Proposition 3

If a viable late deal exists, then by Proposition 1, the no-deal equilibrium of the subgame where A knows B's program is second-stage must feature immediate costly conflict—if it did not, then no equilibrium of this subgame would entail costly conflict. This implies that there is a positive surplus to be gained from agreeing to a late deal rather than playing this no-deal equilibrium. For a late deal to be viable, neither player can receive more than the entire surplus under the deal; otherwise, at least one player will have a profitable deviation to reneging on the deal by initiating costly conflict. Ignoring the knife-edge case in which the viability of a late deal requires that the entire surplus from the deal must be given to B, this implies that there is a viable late deal which renders both players strictly better off than the no-deal equilibrium. Any viable early deal must be supported by the threat of reverting to the no-deal equilibrium if either player reneges on the deal, since only the no-deal equilibrium features a positive probability of costly conflict. But then, no such deal can be renegotiation-proof, as if either player reneges, the two will with positive probability arrive at a point where costly conflict is called for, but at which both would do strictly better to agree to a late deal instead.

Next we demonstrate that, as long as ϵ is low enough, any viable late deal must be more generous to B than any viable early deal. Any viable early deal yields a continuation value for $B1$ of no more than $\overline{V_{d1}^B} = 1/(1-\delta) - V_{nd,0}^A$, where $V_{nd,0}^A$ is the lowest value A could receive from a no-deal equilibrium of a subgame where it is common knowledge that B's program is in the first stage. (The zero in the subscripts refers to the number of null stage signals A has received since he was last certain that B's program was in the first stage.) Any deal that was more generous to B would obviously be reneged upon by A. The method of proof is to show that, by choosing ϵ small enough, $V_{nd,0}^A$ can be made arbitrarily close to $1/(1-\delta) - W^B$, so that the value of the best possible deal for $B1$ can be made arbitrarily close to W^B. This suffices to prove the result, since as will be shown in the proof of Proposition 5, the viable late deal that is least generous to B gives a value that does not depend on ϵ and strictly exceeds W^B.

By the proof of Proposition IV, any no-deal equilibrium is characterized by how long A waits to initiate costly conflict without a second-stage signal. So first suppose that the worst no-deal equilibrium for A features A attacking/abandoning after $k < \infty$ consecutive null signals. At the point at which A has received k consecutive null signals, equilibrium requires that A's costly conflict value at least equal the value A would receive from instead making an acceptable offer q to B:

$$\begin{aligned} W^A \geq q + \delta\rho_k \left[\epsilon\lambda V_n^A + \epsilon(1-\lambda)W^A + (1-\epsilon)\left[\sigma V_{nd,0}^A + (1-\sigma)W^A\right]\right] \\ + \delta(1-\rho_k)\left[\lambda V_n^A + (1-\lambda)W^A\right] \end{aligned}$$

As argued in the proof of Proposition II, $q = p + c_B$ is always acceptable to B, and equilibrium requires that $V^A_{nd,0} \geq W^A$, so we have:

$$W^A \geq p + c_B + \delta\rho_k \left[\epsilon\lambda V^A_n + \epsilon(1-\lambda)W^A + (1-\epsilon)W^A\right] \\ + \delta(1-\rho_k)\left[\lambda V^A_n + (1-\lambda)W^A\right]$$

Observe that by choosing ϵ small enough, we can make the first two terms inside the brackets on the first line arbitrarily close to zero, and Proposition III implies we can do the same for the term on the second line. Thus, for sufficiently small ϵ, we can rewrite the inequality as:

$$W^A \geq p + c_B + \delta W^A + \text{arbitrarily small terms} \\ \Leftrightarrow p - c_A \geq p + c_B + \text{arbitrarily small terms}$$

This contradiction implies that, for ϵ sufficiently low, it must be that $k = \infty$. That is, A will not attack/abandon in the no-deal equilibrium, no matter how many consecutive null stage signals are received, until a second-stage signal occurs.

So now suppose that in the worst no-deal equilibrium for A, A will never initiate costly conflict without a second-stage signal. After receiving i consecutive null stage signals, A's continuation value in the no-deal equilibrium is given by:

$$V^A_{nd,i} = q_i + \delta\rho_i \left[\epsilon\lambda V^A_n + \epsilon(1-\lambda)\left[\sigma W^A + (1-\sigma)V^A_{nd,i+1}\right] + (1-\epsilon)\left[\sigma V^A_{nd,0} + (1-\sigma)V^A_{nd,i+1}\right]\right] \\ + \delta(1-\rho_i)\left[\lambda V^A_n + (1-\lambda)\left[\sigma W^A + (1-\sigma)V^A_{nd,i+1}\right]\right]$$

It was shown in the proof of Proposition IV that A's continuation value in any no-deal equilibrium is strictly declining in the number of consecutive null stage signals A has received. Equilibrium requires that these values be bounded below by W^A, so we know that $\lim_{i\to\infty} V^A_{nd,i} \equiv V^A_{nd,\infty}$ exists and that, by choosing a high enough i, $V^A_{nd,i}$ can be made arbitrarily close to $V^A_{nd,i+1}$. Using again the fact that $q_i > p + c_B$, the equation above, for sufficiently high i, implies:

$$V^A_{nd,i} > p + c_B + \delta\rho_i \left[\epsilon\lambda V^A_n + \epsilon(1-\lambda)\left[\sigma W^A + (1-\sigma)V^A_{nd,i}\right] + (1-\epsilon)\left[\sigma V^A_{nd,0} + (1-\sigma)V^A_{nd,i}\right]\right]$$
$$+ \delta(1-\rho_i)\left[\lambda V^A_n + (1-\lambda)\left[\sigma W^A + (1-\sigma)V^A_{nd,i}\right]\right] - \text{arbitrarily small terms}$$

Similarly to the above, by choosing ϵ small enough we can shrink many of the terms arbitrarily close to zero, and using $V^A_{nd,0} \geq V^A_{nd,i}$, we arrive at:

$$V^A_{nd,i} > p + c_B + \delta V^A_{nd,i} - \text{arbitrarily small terms}$$
$$\Leftrightarrow V^A_{nd,i} > \frac{1}{1-\delta} - W^B - \text{arbitrarily small terms}$$

Equilibrium requires that $V^A_{nd,i} \leq 1/(1-\delta) - W^B$, since otherwise B would reject A's offer in favor of costly conflict. It follows that by choosing i high enough, and ϵ low enough, we can make $V^A_{nd,i}$ arbitrarily close to $1/(1-\delta) - W^B$. Since $V^A_{nd,0} \geq V^A_{nd,i}$, the same is true for $V^A_{nd,0}$. □

4 Conditions for Deal Viability

We have analyzed whether a deal could ever be agreed and when the most propitious moments for a deal occur. Here, we discuss how the exogenous parameters of the model affect whether a deal is actually viable at either of these moments. Knowledge of these effects can inform policymakers' attempts to manipulate these parameters in pursuit of nonproliferation. Unfortunately, these effects are often surprising, and sometimes hard to predict. For clarity, we first present the highlights here. A comprehensive analysis including formal propositions and proofs follows.

In particular, we are interested in the effects on the viability of late and early deals of the proliferation-induced shift in power $(p - p_n)$, the costs of attacking/abandoning to prevent it ($c_A + c_B$), the technological sophistication of the proliferant's program (ϵ and λ), and the quality of A's monitoring of the program (τ and σ). All of these vary across cases, and all are potentially policy-manipulable.

We begin with a late deal, occurring when A believes B's program is nearing success and enforced by the threat of immediate resort to costly conflict. As one might expect, the better A's monitoring of B's compliance is (higher τ), the more likely it is that a deal will be viable, as B's cheating is more likely to be caught and punished. More interestingly, the effects of the other parameters are non-monotonic. If the effect of proliferation ($p - p_n$) is too low, the costs of conflict ($c_A + c_B$) are too high, or the sophistication of B's program (λ) is too low, then B's program is not dangerous enough to A to be worth costly conflict, and so A's threat is not credible and no deal is enforceable. For higher effects of proliferation, lower costs, or higher sophistication, A's threat becomes credible and a deal is viable. However, if the effect of proliferation gets much higher, conflict costs much lower, or sophistication much higher, A's threat remains credible, but the margin by which the expected penalty exceeds the temptation to cheat will decline, and eventually the temptation to cheat will overwhelm the penalty and the deal will again be unenforceable. Intuitively, the higher the effects of proliferation or the sophistication of B's program, the more tempting it is for B to cheat on an agreement. The lower the costs of conflict, the less willing A will be to offer the concessions necessary to induce B's compliance. Thus proliferation effects, conflict costs, and program sophistication have competing consequences for a deal's viability. A deal exists only when the values of these variables are in a "sweet spot," neither too high nor too low.

This implies that policies aimed at mitigating the problems of proliferation may sometimes have the opposite of the intended effects. Clearly, efforts on the part of the United States to improve the quality of monitoring embedded nonproliferation agreements, as well as its own intelligence capabilities, unambiguously raise the enforceability of late deals. Lowering the costs of preventive attack, say by developing improved capabilities for destroying deeply-buried nuclear facilities, can strengthen the viability of agreements if it makes the US willing to attack where it previously wasn't. However, it can instead *undermine* agreements if it simply decreases the cost of a conflict the US was already willing to wage in the absence of a deal. Efforts to lessen the effects of proliferation, such as the development of missile defenses and preemptive strike capabilities, may render a deal viable by making a proliferant willing to abide by it, since the value of having nuclear weapons is lower. However, if these effects are modest to begin with, lessening them can undermine a deal if the US becomes unwilling to attack to stop the proliferant in the absence of a deal. Finally, raising the difficulty of mastering the latter stages of weapons development, such as with sabotage or international restrictions on the transfer of relevant materials or expertise, has similarly context-dependent effects. The bottom line is that the effects of these policies can only be assessed on a case-by-case basis—they are not necessarily always helpful, or even helpful on average.

What about an early deal, agreed when A knows B's program to be rudimentary? As with a late deal, the likelihood of an early deal's viability increases in the quality of her monitoring of his program (τ and σ). Better monitoring makes the relatively good possible outcomes for B of cheating, such as making it to a later stage and continuing to cheat under A's nose, less likely, and the relatively bad outcomes, such as costly conflict or the end of the deal the concessions B received under it, more likely. It also lowers the value of many of the cheating possibilities, because it reduces the expected time until B gets caught cheating or at a late stage of progress, leading to attack or abandonment. It also generally raises the surplus—higher σ generally means a lower value in the no-deal equilibrium because of the higher likelihood of costly conflict when A catches B's program at a late stage of progress—so that A has more to offer B in exchange for compliance.

The effects of changes in the costs of conflict, the effect of proliferation, and the sophistication of B's program on the viability of an early deal are much more complicated.

For instance, an increase in the sophistication of B's program (ϵ and λ) increases the likelihood that cheating leads to the better outcomes for B—rapid progress and success in his program—raising his temptation to cheat. But it also generally increases the surplus available to support a deal, because it leads A to resort to conflict sooner in the no-deal equilibrium due to worries of rapid, unobserved progress in B's program. The increased surplus enables larger concessions to B to support a deal, making it more likely to be viable. Which of these countervailing effects dominates depends in a complicated way on the values of the other parameters. Even more so than with a late deal, the implication is that it is very difficult to predict *ex ante* whether a proposed policy to shift one of these parameters will contribute to or undermine the viability of any particular early deal, or early deals in general.

To complete the analysis, we present the formal conditions under which late and early deals are in equilibrium.

Proposition 5. *There is a viable late deal if and only if* $p - c_A + \frac{\delta\lambda}{1-\delta}(p - p_n) > 1 + \frac{\delta\lambda}{1-\delta}(c_A + c_B)$ *and* $[\lambda + \tau(1-\lambda)](c_A + c_B) \geq \lambda(p - p_n)$.

Intuition:

The first condition is actually equivalent to Assumption 1: it guarantees that once A is certain B's program has reached the advanced stage of development, then if no deal is agreed, costly conflict will occur. This means that there is potentially something for *both* sides to gain from agreeing on nonproliferation: they will not have to pay the costs of conflict. It also means that A has a credible threat with which to try to enforce a deal, and so all that matters is whether that threat is severe enough to overcome B's temptation to cheat on the agreement. The second condition formalizes this requirement. If A is to get at least as much value from a deal as from the no-deal equilibrium (i.e., costly conflict), then the most she could possibly give B to encourage compliance with the deal, and thus the most she could threaten to take away if he is caught cheating, is the surplus value that would be destroyed by the conflict ($c_A + c_B$, per period). Since A can only punish B if his investment is either detected or successful (in which case A knows B has cheated because he has nuclear weapons), this penalty is weighted by the probability that either occurs ($\lambda + \tau(1-\lambda)$). For B, the temptation to cheat is the shift in the balance of power that proliferation would effect, weighted by the chance that his program would succeed (or $\lambda(p - p_n)$). If the maximum expected penalty exceeds the expected benefit of cheating, then the deal is viable; otherwise, one or the other side will not comply with any postulated agreement.

Proof. We will prove this result in two steps. First, assume that there is a deal equilibrium. Then, according to Proposition 1, the no-deal equilibrium of the subgame must involve costly conflict, and by Proposition I, the first condition has to be true. Then, under the deal, B must prefer not investing to investing. More specifically, it must be that

$$V_d^B \geq \lambda V_n^B + (1-\lambda)(\tau W_B + (1-\tau)V_d^B) \tag{2}$$

This is easiest to satisfy when B's value under the deal is maximized. The most value A can yield to B is $\frac{1}{1-\delta} - W_A$, which makes A indifferent between making the corresponding offers and initiating costly conflict. Substituting $V_d^B = \frac{1}{1-\delta} - W_A$ in the above inequality and rearranging terms gives the second condition.

For the second part of the proof, assume that the two conditions hold. Consider the following strategy profile as a candidate for a deal equilibrium. So long as the deal holds—i.e., A has always offered $q \leq p - c_A$ and B has always accepted this and never been observed to invest—A offers $q = p - c_A$ and B accepts this or any more generous offer and does not invest. If A makes a stingier offer or B is observed to have invested, then until costly conflict or proliferation occurs, A attacks/abandons at the first opportunity and B accepts any offer and invests.

The first condition specified in the proposition guarantees that, once the deal no longer holds, the strategies of the players form an equilibrium, as under Proposition I. Given B's strategy, as long as the deal holds, A cannot profitably deviate by initiating costly conflict, and hence also cannot profitably deviate to making a stingier offer, since he would prefer costly conflict to this. A also cannot profitably make a more generous offer to B, since this lowers A's payoff in this period and leaves his future value unchanged. Given A's strategy, as long as the deal holds, B cannot profitably deviate to rejecting an offer of $q \leq p - c_A$, since doing so yields B's costly conflict value while accepting such an offer yields at least $\frac{1}{1-\delta} - W^A > W^B$. Because the second condition specified in the proposition is equivalent to inequality 2 above with $V_d^B = \frac{1}{1-\delta} - W^A$, B also cannot profitably deviate to investing. □

Proposition 6. *There is a viable early deal if and only if:*

$$\frac{1}{1-\delta} - V_{nd}^A \geq \epsilon\lambda V_n^B + \epsilon(1-\lambda)\left[\sigma W^B + (1-\sigma)\tau V_{nd,1}^{B,s2} + (1-\sigma)(1-\tau)V_{cheat}^{B,s2}\right] + (1-\epsilon)\left[\sigma\tau V_{nd}^{B,s1} + (1-\sigma)\tau V_{nd,1}^{B,s1} + (1-\tau)\left(\frac{1}{1-\delta} - V_{nd}^A\right)\right]$$

Intuition:

The left-hand side of the condition is the highest value B could receive in any deal that A would be willing to offer: the total value of the game $(1/1-\delta)$, less the value A would get in the no-deal equilibrium (V_{nd}^A). A can't offer any more than this and still prefer the deal to its absence. The right-hand side is what B would get by cheating on the putative deal. If B's investment leads immediately to nuclear weapons, then he gets the greater concessions from A that proliferation brings (V_n^B). If his investment leads to mastery of the first stage but not the second, what happens depends on what A observes. If she detects that B has reached the second stage, then she will resort to conflict immediately (giving B the value W^B), but if she detects only that he has cheated, then the no-deal equilibrium will begin with him in the second stage and her having seen a period of investment but unsure of his stage (denoted $V_{nd,1}^{B,s2}$). If A detects nothing, then B has the opportunity to continue cheating on the deal without her knowing ($V_{cheat}^{B,s2}$). And finally, if B's investment goes nowhere, and A sees both stage and investment, then the no-deal equilibrium begins (giving $V_{nd}^{B,s1}$). If instead A detects B's cheating but not his stage, then the no-deal equilibrium begins with him in the first stage and her having gone one period without a stage signal (yielding $V_{nd,1}^{B,s1}$), and if she does not see his cheating, then B can go back to compliance with the deal with A none the wiser. If the left-hand side exceeds the right, then there is a deal A is willing to offer that is generous enough to induce B's compliance.

Proof. First, assume there is a deal equilibrium in a subgame in which B's program is known to be in the first stage. Observe that the most tempting opportunity for B to renege on the deal by investing will be in the first period of the deal. B's possible continuation values from initiating investment are identical in different periods, except for the continuation value that results when, in the period after B reneges, A observes B's investment but not the program's stage. Though investment is off the equilibrium path, so that A's belief after observing investment is unconstrained by weak consistency, the only empirically plausible belief A could hold immediately after observing that B had invested in the first period of the deal is that B has invested for exactly one period. By contrast, having first observed investment in any later period, A might reasonably believe that B had been investing for multiple periods but only the last was detected. B's value under a no-deal equilibrium is declining in the number of periods he is believed to have invested, since this means that the time at which A would resort to costly conflict rather than tolerate further investment is nearer, and in the meantime B is only receiving positive offers if these are necessary for A to minimally satisfy B and avoid costly conflict. It follows that the most tempting time to invest is at the beginning of the deal.

Thus, for the deal to be in equilibrium, it must be that:

$$\begin{aligned} V_d^B &\geq \epsilon\lambda V_n^B + \epsilon(1-\lambda)\left[\sigma W^B + (1-\sigma)\tau V_{nd,1}^{B2} + (1-\sigma)(1-\tau)V_{cheat}^{B2}\right] \\ &+ (1-\epsilon)\left[\sigma\tau V_{nd,0}^{B1} + (1-\sigma)\tau V_{nd,1}^{B1} + (1-\tau)V_d^B\right] \end{aligned} \tag{3}$$

The most A can offer to B in any deal is $\frac{1}{1-\delta} - V_{nd,0}^A$; otherwise, A prefers the no-deal equilibrium. Substituting this term for V_d^B gives the condition in the proposition.

Second, we prove the converse. Assume that the condition given in the proposition holds, and that immediate costly conflict is the equilibrium outcome when B's program is observed to be in the second stage. Then by Propositions II and IV, in any no-deal equilibrium of the game, B always invests, and A tolerates B's program until he becomes sufficiently confident that it has reached the second stage, and then attacks/abandons. Denote A's and B's values under such an equilibrium V_{nd}^A and V_{nd}^B.

Consider the following strategy profile as a candidate for a deal equilibrium. So long as the deal holds—i.e., A has always offered $q \leq (1-\delta)V_{nd}^A$, B has always accepted these offers and has never been observed to invest or to have reached the second stage—A offers $q = (1-\delta)V_{nd}^A$ and $B1$ accepts this or any more generous offer and does not invest. If A makes a stingier offer or B is observed to have invested or to have reached the second stage, then the two players immediately revert to the no-deal equilibrium corresponding to A's belief at that point about the stage of B's program. Regardless of whether the deal holds or not, $B2$ always accepts any offer and invests.

By construction, once the deal no longer holds, the players' strategies form an equilibrium. The proof of Proposition II shows that, given A's strategy, $B2$'s is a best response. Given $B1$'s strategy, as long as the deal holds, A cannot profitably deviate to attacking/abandoning, since $V_{nd}^A \geq W^A$. A also cannot profitably deviate to making a stingier offer, since this would induce $B1$ either to reject the offer in favor of costly conflict (if it were low enough) or to invest, and either would yield a value no more than V^A_{nd}. A cannot do better with a more generous offer to B, since this lowers A's payoff in this period and leaves his future value unchanged. Given A's strategy, as long as the deal holds, $B1$ cannot profitably deviate to rejecting an offer of $q \leq V^A_{nd}(1-\delta)$, since doing so yields B's costly conflict value while accepting it gives B at least $\frac{1}{1-\delta} - V^A_{nd} \geq W^B$. Because the condition given in the proposition is equivalent to inequality 3 above, $B1$ also cannot profitably deviate to investing. □

5 Additional Empirical Cases for Hypotheses 2 and 3

H2: Late Deals and Threat Credibility

The December 2003 deal with Libya was also attended and catalyzed by a newly-credible US threat to attack Libya in the absence of a deal. Although the US and Libya engaged in serious conflict in the 1980s, this was related to Libya's involvement in international terrorism, not its nuclear program: the US never attacked Libya's nuclear facilities or even seriously considered doing so up to 2000.[3] By contrast, during the buildup to the Iraq War, the US issued threats to Libya insinuating that, because it had behaved similarly to Iraq, it might receive the same treatment, and included Libya on a classified (but leaked) list of states that might be subject to preventive attack due to pursuing WMD.[4] Gaddafi apparently viewed these threats as credible, and cited his desire to avoid Saddam's fate as motivation for the deal.[5]

The US and South Korea seriously considered launching preventive strikes against North Korea's nuclear facilities and publicly threatened to do so in the prelude to the October 1994 Agreed Framework. Though South Korea had considered this option previously in 1991 and 1993, it would not have undertaken strikes unilaterally, and the US did not begin planning

[3] Fuhrmann and Kreps (2010).
[4] Corera (2006, 182) and Jentleson and Whytock (2005/06, 73).
[5] Bowen (2006, 64).

an attack until late 1993.[6]

Taiwan began to pursue the capability to produce nuclear weapons in the late 60s, leading the US to negotiate with it to end its efforts. Taiwan made major concessions on three occasions: January 1977, when it ceased attempts to purchase a reprocessing facility from abroad; the spring of 1977, when it allowed the US to directly oversee its nuclear research and cancelled those lines of work the US considered dangerous; and September 1978, when it ceased enrichment research.[7] Each of these concessions followed on the heels of a US threat that continued activity might lead the US to abandon Taiwan, and such threats were not issued at any previous point.

Finally, the US negotiated with Ukraine to eliminate the nuclear weapons and associated infrastructure Ukraine inherited from the defunct USSR in 1991.[8] Ukraine gave up these weapons in three steps: from December 1991 to May 1992, it allowed the transfer of its tactical nuclear weapons to Russia; from July 1993, when it began the process of dismantling its SS-19 nuclear missiles; and from March 1994 to June 1996, when the rest of its nuclear weapons were removed.[9] Each of these steps was directly preceded by a US threat, and such threats were not issued at any other point.[10]

[6] See the appendix to Fuhrmann and Kreps (2010), pp. 7–8 and 13–14.

[7] Our chronology of concessions and threats is drawn from Burr (2007).

[8] Though Ukraine and the US were not allies, in the aftermath of the Soviet collapse, Ukraine believed that the US had the power to avert the danger Russia posed to Ukraine's territorial and economic integrity (Budjeryn, 2016, 110–112, 126–131, 137). Thus the relevant US threat was to abandon Ukraine to Russian coercion and possibly attack.

[9] Budjeryn (2016, 94, 114, 134 ,168, 178).

[10] The first step followed the US explicitly conditioning diplomatic recognition of Ukraine's independence on its denuclearization (Budjeryn, 2016, 61–62). The exact US threats that preceded the second and third steps are not discernible from available sources, but they were made directly by US Secretary of State Baker and President Clinton to Ukrainian Foreign Minister Zlenko and President Kravchuk, and were described by observing US diplomats as extremely harsh (Budjeryn, 2016, 137–138, 175–176). We therefore assume they were at least implicitly threats of abandonment.

H3: Late vs. Early Deals

Iran got a European guarantee against US attack and removed the possibility of multilateral sanctions as long as the agreement lasted—in exchange only for Iran's agreement to temporarily suspend uranium enrichment, something that slowed but did not stop its program's progress.[11]

North Korea was promised the construction of two advanced nuclear reactors and enormous shipments of heavy oil until the reactors were completed, all free of charge, as well as movement toward normalization of relations with the US.[12]

[11] Volpe (2015, 217–219).
[12] International Atomic Energy Agency (1994).

References

Bowen, Wyn Q. 2006. "Libya and Nuclear Proliferation: Stepping Back from the Brink." *The Adelphi Papers* 46(380):7–103.

Budjeryn, Mariana. 2016. The Power of the NPT: International Norms and Nuclear Disarmament of Belarus, Kazakhstan, and Ukraine, 1990–1994 PhD thesis Central European University.

Burr, William. 2007. U. S. Opposed Taiwanese Bomb during 1970s. Electronic Briefing Book 221 National Security Archive.

Corera, Gordon. 2006. *Shopping for Bombs: Nuclear Proliferation, Global Insecurity, and the Rise and Fall of the AQ Khan Network*. Oxford University Press.

Debs, Alexandre and Nuno P. Monteiro. 2014. "Known Unknowns: Power Shifts, Uncertainty, and War." *International Organization* 68(1):1–31.

Fuhrmann, Matthew and Sarah E. Kreps. 2010. "Targeting Nuclear Programs in War and Peace: A Quantitative Empirical Analysis." *Journal of Conflict Resolution* 54(6):831–859.

International Atomic Energy Agency. 1994. "Agreed Framework of 21 October 1994 between the United States of America and the Democratic People's Republic of Korea." INFCIRC/457.

Jentleson, Bruce W. and Christopher A. Whytock. 2005/06. "Who 'Won' Libya? The Force-Diplomacy Debate and Its Implications for Theory and Policy." *International Security* 30(3):47–86.

Volpe, Tristan A. 2015. Proliferation Persuasion: Coercive Bargaining with Nuclear Technology PhD thesis The George Washington University.