A Dynamic Theory of Nuclear Proliferation and Preventive War
Online Appendix

Muhammet A. Bas and Andrew J. Coe

A1 Proofs of Propositions

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

Lemma 1. *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 war will occur.*[1]

Subgame perfection implies that B will accept any $q < p_n + c_B$, because rejecting it yields his war value, while accepting it and going to war 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 war. 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 war. Thus, in equilibrium A makes the offer that renders B indifferent to war ($q = p_n + c_B$) and B accepts this or any higher offer and rejects any lower offer. □

Proposition 1

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] 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.

$1/(1-\delta) - V_n^B = (p_n + c_B)/(1-\delta)$. 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.

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 starting a war, 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 war value $W^A \equiv (p - c_A)/(1-\delta)$, then the unique no-deal equilibrium features immediate war. 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 war condition given in the statement of the proposition.

Suppose instead that $1/(1-\delta) - \underline{V}^B > W^A$ and that war occurs in equilibrium. In the period in which war 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 going to war in the next round yields a better value for B than starting a war 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 war, so at least one player has a profitable deviation to not starting a war, and thus war 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. $\square$

Proposition 2

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 war value W^B. If it is not, B could profitably deviate by rejecting A's offer and thereby causing a war. 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 war value. If the investment fails, he will receive at least his war value in the next round, but it will succeed (in the sense of B getting nuclear weapons) with positive probability and, by Lemma 1, 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 war value. Thus, in a no-deal equilibrium, A's offers 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, and if she 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 1 applies, and

$B2$ will always invest, given the chance.

Moreover, the supposition that a second-stage signal leads to war implies that $B2$ strictly prefers investment to war, 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 1 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 war, 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 rather than making the offer; otherwise, the offer makes war certain and so we treat it as equivalent to attacking.

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:

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

$$V_t^{B1}(\mathrm{NI}) = 1 - q_t + \delta\left[\sigma V_0^B + (1-\sigma)V_{t+1}^{B1}\right]$$

$$V_t^{B2} = 1 - q_t + \delta\left[\lambda V_n^B + (1-\lambda)\left(\sigma W^B + (1-\sigma)V_{t+1}^{B2}\right)\right]$$

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 $V_t^{B1}(\mathrm{I}) \leq V_t^{B1}(\mathrm{NI})$, 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 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 B unless she 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 war, which exceed those associated with A's pre-proliferation offers, and the formula ignores discounting due to any delay in war 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 war 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 3

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 4

Proposition 2 establishes B's optimal behavior in a no-deal equilibrium, and Proposition 3 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 2 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 war and peace, or B will accept any offer regardless of stage. If there is an offer that renders $B1$ indifferent between war and peace, it cannot be in equilibrium for A to offer more: by Proposition 2, 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 2 that it cannot be in equilibrium for A to offer less (A would strictly prefer to attack). 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 war or proliferation, and up to the occurrence of a first-stage signal that would reset A's estimate of the probability she 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 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 attacking.

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 her estimate of the probability she faces $B2$ is high enough to merit attacking rather than tolerating further risk of proliferation. It also implies that once A has reached this threshold, she will attack 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 attack 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 3, as additional consecutive null signals are received, A's estimate of the probability that she faces $B2$ converges, so that any strategy excluded by this restriction would require A to randomize over attacking at some point arbitrarily close to her 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 in that period or after any further consecutive null signals. (Not attacking 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 attacking.)

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 attacking after i consecutive null signals.

$$V_i^{B2} = 1 - q_i + \delta\left[\lambda V_n^B + (1-\lambda)\left[\sigma W^B + (1-\sigma)V_{i+1}^{B2}\right]\right]$$

$$V_i^{B1} = 1 - q_i + \delta\epsilon\left[\lambda V_n^B + (1-\lambda)\left[\sigma W^B + (1-\sigma)V_{i+1}^{B2}\right]\right] + \delta(1-\epsilon)\left[\sigma V_0^B + (1-\sigma)V_{i+1}^{B1}\right]$$

$$V_i^A = q_i + \delta\left[\rho_i\epsilon + (1-\rho_i)\right]\lambda V_n^A + \delta\left[\rho_i\epsilon + (1-\rho_i)\right](1-\lambda)\sigma W^A$$

$$+ \delta\rho_i(1-\epsilon)\sigma V_0^A + \delta\left[\rho_i(1-\epsilon) + \rho_i\epsilon(1-\lambda) + (1-\rho_i)(1-\lambda)\right](1-\sigma)\tilde{V}_{i+1}^A$$

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 she 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 attack 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 war value after $i+1$ consecutive null signals with certainty, but will receive at least his war 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 attacking is in equilibrium only if both attacking 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 3, $\rho_{i-1} > \rho_i$. 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 war in the absence of a non-null stage signal or proliferation, and since every additional consecutive null signal brings a probability of war at least $0 = \pi_{j-1}$, the overall probability that the game will end in war 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 at the first opportunity, then A would not attack no matter how many additional consecutive null signals she 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 war 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 war value. In either case, A could profitably deviate by changing her offer to q^* or attacking. 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 war 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, war, 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 if she knew she 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 3, 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. □

A2 Empirical Evidence

We follow Montgomery and Mount (2014) (henceforth MM) in taking India's program to have acquired nuclear weapons capability by the time of its "peaceful nuclear explosive" in 1974, and in dividing Iran's program into pre- and post-revolution episodes and Iraq's into pre- and post-Gulf War episodes, since in both cases the program effectively had to be restarted from the first stage. For the same reason, we also divide Iran's program at the last of Iraq's successful strikes against the Bushehr reactor in 1988, Iraq's program at Israel's successful strike on the Osiraq reactor in 1981, and North Korea's program at the 1994 Agreed Framework.

We update the preventive attacks data from Fuhrmann and Kreps (2010) (henceforth FK), which end at the year 2000, by adding the 2003 US invasion of Iraq and Israel's 2007 strike on Syria's nuclear program. We also make several changes to the data. First, we stipulate that in the absence of the Gulf War, Iraq would have acquired nuclear weapons. This reflects the conventional wisdom about what would have occurred were it not for the Gulf War, which was not caused by Iraq's nuclear program and can thus be taken as an exogenous interruption of the proliferation-prevention interaction. Next, we assume that in the cases of Libya in 2003, the United States seriously considered preventive attack, and Iran in 2003 and 2005, the United States and/or Israel seriously considered preventive attack. These events occurred after the FK data ends (at 2000), and can only be tentatively imputed given the lack of declassified primary sources on these cases.[2] Finally, we also drop Pakistan's SCoA against India in 1984, and the US/UK attacks against Iraq in 1993 and 1998, on the

[2]Our test results are qualitatively unaffected if these changes are dropped.

grounds that all three were primarily intended for retaliation or punishment rather than for prevention, and thus are not appropriate for testing our theory. According to the appendix of FK, Pakistan considered attacking only in retaliation against an anticipated attack by India against Pakistan's nuclear facilities, rather than to prevent India acquiring the bomb (which it had already mastered in 1974) (A7). The US and UK attacked Iraq in order to punish it for not complying with the settlement imposed on it after the Gulf War. The nuclear facilities struck were known to be inactive and under inspectors' seal, and the intent of the strikes was coercive, not preventive (Devroy and Gellman, 1993; Pollack, 2002, Ch. 3).

We include additional data for the nine episodes not covered by the compilation of intelligence estimates in MM: the US, Japan, the UK, Australia, Egypt, Iran's program under the Shah and during the Iran-Iraq War, Iraq's program before the Osirak strike, and Syria. We do not need estimates of the US or UK programs, since there was no viable potential attacker in both these cases, as argued in the main body on page 25. In the cases of Australia (Walsh, 1997), Egypt (Walsh, 2001), Iraq 1973–81 (Braut-Hegghammer, 2011), Iran 1974–78 and 1984–88 (Koch and Wolf, 1997), and Syria 2001–07 (Albright and Brannan, 2008), we assume no potential attacker estimated these programs to be nearing success, since all were clearly in the first stage, having not yet even completed construction of a suitable reactor, the sole path to fissile material these states were pursuing. For Japan 1943–45, the most likely potential attacker is the US, which assessed Japan as lacking the industrial capacity to develop nuclear weapons and so ignored its program (Grunden, 1998). Finally, for Iran after 2002 (the last estimate in MM), sources suggest that the US and Israel estimated Iran's program to be within four years of success from 2003 to early 2005 (Nuclear Threat Initiative, 2011; Corera, 2006, Ch. 3). However, in 2005, new developments apparently led the US to push the time Iran was expected to get nuclear weapons back to at least 2010 (National Intelligence Council, 2007, 9). Additionally, the appendix to FK implies that Israel would have had a near-success estimate of Pakistan's program in 1979, due to intelligence shared by the US (A5). We also assume that the USSR had a near-success estimate of South Africa's program after 1976, since unlike the US, the Soviets were tracking the program closely, and detected South Africa's construction of a test site (A8). Last, we read the account of US intelligence on North Korea's program from 1995 to 2006 given in the appendix to MM as implying that the US estimated, by 2001, that North Korea had already obtained nuclear weapons sometime in the mid-1990s, and thus never estimated its program to be nearing success in the 1995 to 2006 episode.[3]

A2.1 Measures for Hypothesis 2

For the 22 episodes included in the test, we first need to establish when the likely potential attackers started closely watching the program in question. Unfortunately, we do not have intelligence estimates from any potential attacker except the US. We take the conservative approach of setting the first year of intelligence monitoring for each state to be either the first time serious consideration of attack occurs or the year of the first US estimate that appears in MM, whichever is earlier. In many cases, intelligence is shared between the US and other potential attackers (e.g. Israel, South Korea, UK), so that our measured year will be accurate for these other attackers as well. Other cases mainly involve states toward which the US is less hostile, and MM finds that the US generally pays less attention to the programs of states toward which it is more friendly. This suggests that the likely potential attackers in these cases would start monitoring the programs at least as early as the US, but including earlier years would only strengthen our results.

For the seven episodes included in this test that MM does not cover, we measured the

[3]Given the uncertainty surrounding this episode, one might be inclined to exclude it from the test. Similarly, it could be argued that in the cases of Australia, Egypt, South Korea, Iran 1974–78, Argentina, and Brazil, no potential attacker would ever find it worthwhile to attack these programs, and they should thus also be excluded. Dropping these episodes does not qualitatively alter our results.

beginning of intelligence monitoring as follows. US monitoring of Japan must have begun by 1945, since the US at some point during the war assessed Japan's industrial capacity to pursue the bomb (Grunden, 1998). US monitoring of Australia began by no later than 1968, when the US sent an ACDA/AEC team to investigate Australia's hesitance to sign the NPT (Walsh, 1997). Monitoring of Egypt began even earlier than Egypt's program, with an NIE on its program in 1963 (Central Intelligence Agency, 1963). The US engaged in nuclear cooperation with Iran under the Shah, and from the beginning considered the possibility of a weapons program (Burr, 2009). Sadot (2015) dates the beginning of Israel's careful monitoring of Iraq's nuclear program to 1974. Iran's restarted program was attacked in its first year of existence, so we assume Iraq's intelligence monitoring of it had begun by then. Finally, while there are reports that Israel began monitoring Syria's nuclear program in 2004 (Thomas, 2015), we have not been able to confirm these elsewhere. We instead date the start of monitoring from the US discovery of the site of Syria's reactor, in 2005 (Albright and Brannan, 2008).

Similarly, we must measure when the end of any opportunity for preventive attack occurred, a non-trivial task given that the year when nuclear weapons were actually acquired is not known with certainty in some cases (e.g., Israel, Pakistan). We set the end of each episode to be either the last year in which an attack was seriously considered or the lesser of the years of acquisition given in SW and JG, whichever is later.

Next we need to determine when a program was estimated to be "nearing success." For the same reasons as above, we use the US estimates in MM, and define "nearing success" as occurring when the lower bound of the estimate's most likely range of time until the state in question will get nuclear weapons is four years or less.

A3 Comparative Statics: Propositions and Intuitions

We proceed to analyze the effects of the exogenous parameters on expected behavior in the absence of a deal. We assume throughout that the players' discount rate, δ, is relatively high. If lower values of δ are allowed, the statement of the comparative statics becomes more complicated. However, only relatively high values of δ seem plausible empirically.

Table A1 summarizes the results. The rows contain the exogenous parameters, as well as the endogenous parameter k, which is how long A will wait before attacking in the absence of new intelligence that B's program has not advanced to the second stage. Some of the parameters can affect the outcomes in two general ways: directly, by altering the probability of different paths through the game even as equilibrium k remains constant; and indirectly, by changing the equilibrium value of k, which itself affects the expected outcomes. So, each of these parameters has separate listings for its direct effect on each property and its indirect effect through k, in addition to its overall effect. $+$ and $-$ indicate uniformly positive and uniformly negative effects; 0 indicates no effect; $+/-$ indicates that the direction of effect depends on the other parameters; one sign is circled if it tends to predominate.

Lemma 2. *Increasing k leads to a higher probability of proliferation, lower probability of war and mistaken war, and a longer expected time to proliferation or war.*

The longer A is willing to wait before attacking in the absence of new intelligence (k), the more chances B will have for his investment to bear fruit, so that the probability of proliferation increases and the probability of war decreases. Correspondingly, there are more chances for the game to end before A's mounting suspicion leads her to attack in the absence of new intelligence, so that the probability of mistaken war—recall, in which A attacks when B's program has not actually advanced to the second stage—decreases. Finally, because A will wait longer before attacking, the interaction between the two players is also expected to last longer before proliferation or war occurs.

	Endogenous Outcomes				
	k	Pr(Proliferation)	Pr(War)	Pr(Mistaken War)	Expected Time
k		$+$	$-$	$-$	$+$
$p - p_n$	$-$	$-$	$+$	$+$	$-$
$c_A + c_B$	$+$	$+$	$-$	$-$	$+$
λ	$+/\ominus$	$+/\ominus$	$\oplus/-$	$\oplus/-$	$+/\ominus$
direct		$+$	$-$	0	$-$
through k		$+/\ominus$	$\oplus/-$	$\oplus/-$	$+/\ominus$
ϵ	$+/\ominus$	$+/\ominus$	$\oplus/-$	$\oplus/-$	$+/\ominus$
direct		$+$	$-$	$-$	$-$
through k		$+/\ominus$	$\oplus/-$	$\oplus/-$	$+/\ominus$
σ	$\oplus/-$	$+/\ominus$	$\oplus/-$	$+/\ominus$	$+/\ominus$
direct		$+/\ominus$	$\oplus/-$	$-$	$+/\ominus$
through k		$\oplus/-$	$+/\ominus$	$+/\ominus$	$\oplus/-$

Table A1: Comparative Statics

Intuitively, k is governed by a tradeoff. Waiting longer before attacking gives A more time to enjoy the surplus from avoiding war, as well as more time for new intelligence on B's program to come in, possibly revealing that it remains at the first stage, so that A needn't attack after all. But it also exposes A to an increasingly large risk that B's program will succeed, forcing A to offer better concessions once B has nuclear weapons. The equilibrium value of k roughly balances this tradeoff: waiting until later to attack exposes A to more risk than the additional surplus from peace is worth, while attacking earlier eliminates too small a risk to justify the lost surplus. The balance is rough because A can only attack at discrete intervals, so that the decision to attack will not come at the exact point in time at which the value of attacking now and that of delaying one instant more are perfectly equal.[4]

Because the risks of peace and cost of war are not perfectly balanced, small-enough changes in the exogenous parameters will not alter the terms of the tradeoff enough to cause A to prefer to attack an entire period sooner (or later). Thus, small changes in the parameters may have different effects than large changes, as we will see next.

Proposition 5. *A sufficiently large increase in $p - p_n$, or sufficiently large decrease in $c_A + c_B$, will decrease equilibrium k and thereby lower the probability of proliferation, raise the probabilities of war and mistaken war, and shorten the expected time to proliferation or war. Small-enough changes will not affect the likelihood of the outcomes.*

The larger the shift in power $(p - p_n)$ due to B getting nuclear weapons, the greater the risk for A of waiting any longer, and the quicker A resorts to attack. Delaying the attack offers the advantage of putting off its costs $(c_A + c_B)$, but the smaller these are, the less reason there is for A to dally. Thus, if the shift increases enough, or the costs decrease

[4] This occurs because, in our model, time passes in discrete periods rather than continuously, but we do not view this result as an artifact of the model setup. A continuous-time model would entail the implausible assumption that the leadership of state A must continuously reconsider the choice of policy toward the proliferant. It seems more realistic to assume, as our model does, that having decided to tolerate the proliferant's program for now, the government of A will turn to other issues until some time has passed or new intelligence has come in, so that a bureaucratically costly reevaluation of policy toward the proliferant comes to seem justified.

enough, A will attack sooner, so that the probabilities of war and mistaken war increase and the interaction will end sooner. By contrast, if the changes are small enough, the best time to attack will not change, and the likelihood of different outcomes will remain the same.

Proposition 6. *A small-enough increase in λ will increase the probability of proliferation, decrease the probability of war, leave the probability of mistaken war unchanged, and decrease the length of the game, without affecting k. If ϵ is low enough, and λ is close enough to ϵ, then a large-enough increase in λ may increase k, overall raising the probability of proliferation and the expected length of the game, but lowering the probabilities of war and mistaken war. Otherwise, large-enough increases in λ generally decrease k, overall lowering the probability of proliferation and length, but raising the probability of war and mistaken war.*

Although small changes in the ease with which the proliferant can master the second stage and acquire nuclear weapons (λ) do not affect how long A will wait before attacking, they do affect the likelihood of different outcomes directly. Given the same number of chances for B's program to succeed before A attacks, an easier second stage of development raises the probability that proliferation occurs before A attacks, making proliferation more likely, war less likely, and the expected length of the game shorter. By contrast, since mistaken war can only occur if B's program is still in the first stage when A's patience runs out, the ease of mastering the second stage does not affect the probability of mistaken war.

Larger increases in λ have competing indirect effects on how long A is willing to wait. On the one hand, in any given period, it becomes more likely that B's program will advance from its present stage to successfully producing nuclear weapons, raising the risk of proliferation, so that A is encouraged to attack sooner. On the other hand, higher λ means that in any given period, A will estimate that B's program is less likely to have reached the second stage (because if it had, the ease of mastering that stage means that A should have observed B getting nuclear weapons by now). This latter effect tends to encourage A to delay attacking and so reinforces the direct effects of increased λ: increasing the probability of proliferation and expected length of the game, while lowering the probabilities of war and mistaken war. However, this latter indirect effect can only possibly outweigh the former when two conditions are met. First, the change in A's estimate due to increased λ must be large; this occurs when $\lambda \approx \epsilon$. Second, A must be much better off when B's program is in the first stage rather than the second; this requires that ϵ be low, so that the probability of B acquiring nuclear weapons is much lower from the first stage than from the second.

If either of these conditions is not met, then a rise in λ will lead A to attack sooner. The indirect effects of the rise in λ generally outweigh its direct effects, so that the probability of proliferation and expected length of the game go down, while the probabilities of war and mistaken war go up. In other words, A *more than compensates* for the rise in λ by attacking sooner, so that a proliferant that finds it easier to master the second stage of development is counter-intuitively *less* likely to get the weapons. The reason this overcompensation occurs is that, the smaller k is to start, the larger will be the decrease in the probability of proliferation as A reduces k yet further. So, as λ rises and k ratchets down, the decrease in the probability of proliferation due to A attacking sooner exceeds the increase due to rising λ.

Proposition 7. *A small-enough increase in ϵ will increase the probability of proliferation, decrease the probabilities of war and mistaken war, and decrease the length of the game, without affecting k. A large-enough increase generally decreases k and the game's length, overall lowering the probability of proliferation, while raising those of war and mistaken war.*

Just as with λ, a small increase in ϵ will not affect how long A will wait before attacking, but the fact that the first stage of weapons development is easier to master means B is more likely to proliferate before A attacks, so that war is less likely and the expected length of the game declines. Unlike with λ, the increased ease of moving to the second stage means that, if A's patience runs out and A attacks without certainty that B's program has reached the second stage, A is less likely to be mistaken.

Larger increases in ϵ generally reduce A's willingness to wait before attacking. In any given period, it becomes more likely that B will get nuclear weapons, for two reasons: first, because B is more likely to master both stages in one period, since the first stage has become easier; second, because A's estimate that B has already reached the second stage by that period will be higher. The latter reason also implies that the arrival of new intelligence that would allay A's suspicions and make it possible for A to put off the costs of war is also less likely, again encouraging A to attack sooner rather than later.

Similar to λ, the indirect effects through k of a higher ϵ generally outweigh its direct effects, so that the probability of proliferation and expected length of the game counter-intuitively go down, while the probabilities of war and mistaken war go up. This occurs for the same reason: as A decreases the time she is willing to wait before attacking, the probability of proliferation falls more and more as k decreases.

Proposition 8. *Increases in σ will generally increase k and reduce the likelihood of mistaken war. If σ and equilibrium k are low enough, increasing σ may increase the probability of proliferation and the expected length of the game, while reducing the probability of war. Otherwise, increasing σ generally reduces the probability of proliferation and the expected length of the game, while increasing the probability of war.*

Finally, we turn to the effects of changes in A's ability to monitor the progress of B's program (σ). First, an increase in σ generally makes A more willing to wait before attacking in the absence of new intelligence—if she waits, a new signal is more likely to come in. This reduces the risk that A might attack prematurely while B's program remained in the first stage, wasting the surplus from peace she could otherwise safely enjoy, and also the risk of B getting another try at proliferation when his program has already advanced to the second stage without A detecting it. As we saw from Lemma 2 above, this reduces the probability of a mistaken war occurring. But a rise in σ, even if it is too small to change k, will also directly reduce the probability of a mistake. A mistaken war can only occur if A goes long enough without receiving a signal of the stage of B's program, so that A is willing to attack in the absence of definitive intelligence. Increasing σ makes it less likely that A would ever go this long without new intelligence, so that a war launched on the basis of a (possibly erroneous) estimate is less likely to occur.

However, an increase in σ has two opposed direct effects on the likelihood of proliferation and war and the expected length of the game. On the one hand, better monitoring means that A is more likely to catch B when his program is in the second stage, leading to immediate attack. On the other hand, A is also more likely to detect that B's program remains in the first stage, "resetting" A's estimate and so giving B more time to make progress in his program and increasing the chance that the program will succeed. For low enough σ and equilibrium k, the latter effect can dominate: B is likely to be in the first stage for most of the (short) game, so that a (rare) stage signal to A is likely to lead to a reset, giving B more time to succeed. If either of these conditions is not met, then B is more likely to reach the second stage and be exposed to detection and subsequent preventive attack, so that the former effect generally dominates.

When σ's direct effect is opposite to its indirect effect through k, the former generally dominates. While increased σ makes A willing to wait longer before attacking without definitive intelligence, it quickly becomes very unlikely that A would ever go that long without new intelligence being received. Thus, the reduction in the probability of war that comes from lowering the risk of mistaken wars is quickly zeroed out, while the increase in the probability of war that comes from being more likely to catch B once his program has reached the second stage continues to increase. Thus, increasing σ under these conditions will increase the probability of war and decrease the probability of proliferation, while reducing the expected length of the game.

A4 Comparative Statics Proofs

For the remaining proofs, in accordance with the assumption stated above, we take δ to be high (that is, close to 1). We also use numerical simulations of the no-deal equilibria to establish some of our claims: specifically, those that include the modifier "generally" in the statements of the propositions in the main body of the paper. The R code for these simulations is available from the authors on request.

Lemma 2

Proposition 4 establishes that the probability that A will attack in a given period is non-decreasing in the number of consecutive null signals A has received up to the signal (or lack thereof) of that period. So, increasing k to k' can only affect the transition probabilities to war or proliferation in those subgames in which A has received a number of consecutive null signals between k and $k' - 1$. In each of those subgames, where previously A would have attacked with certainty, ending the game, A will now only do so if she receives a second-stage signal; in the absence of such a signal, B will get additional opportunities to invest and possibly succeed in developing nuclear weapons. Thus, the probabilities of war and mistaken war occurring decrease and the probability of proliferation increases in those subgames, and hence overall, and the game lasts longer in expectation. □

Proposition 5

First observe that p, p_n, c_A, and c_B affect the distribution of paths through a no-deal equilibrium only through their effect on k: once equilibrium k is determined, only the parameters λ, ϵ, and σ affect the probability of each possible path. Given the equilibrium value of k, labelled k^*, A must prefer attacking after k^* consecutive null signals to instead making a minimally-satisfactory offer to B (labelled q_k^*) and waiting to attack until the $k^* + 1$ signal has been received, or:

$$W^A \geq q_{k^*} + \delta(1-\rho_{k^*})\left[\lambda V_n^A + (1-\lambda)W^A\right] \\ + \delta\rho_{k^*}\left[\epsilon\lambda V_n^A + \epsilon(1-\lambda)W^A + (1-\epsilon)\left[\sigma V_0^A + (1-\sigma)W^A\right]\right] \quad (2)$$

By Proposition 4, q_{k^*} must be the least generous offer that will satisfy $B1$, or:

$$q_{k^*} = \min\left\{1, 1+\delta\left[\epsilon\lambda V_n^B + (1-\epsilon)\sigma V_0^B + \left[(1-\epsilon)(1-\sigma)+\epsilon(1-\lambda)\right]W^B\right] - W^B\right\}$$

Because k must be a natural number, q_k (through V_0^B), ρ_k, and V_0^A do not change continuously in k, and thus condition 2 will almost never (in the measure-theoretic sense) bind. This implies that small-enough changes in the exogenous parameters, which can all vary continuously, will not cause the inequality to be violated and so will not alter the equilibrium value of k. We use this fact throughout the subsequent proofs.

Similarly, A must prefer attacking after k^* consecutive null signals to instead attacking in the previous period, after k^*-1 signals, or:

$$W^A \leq q_{k^*-1} + \delta(1-\rho_{k^*-1})\left[\lambda V_n^A + (1-\lambda)W^A\right] \\ + \delta\rho_{k^*-1}\left[\epsilon\lambda V_n^A + \epsilon(1-\lambda)W^A + (1-\epsilon)\left[\sigma V_0^A + (1-\sigma)W^A\right]\right]$$

p, p_n, c_A, and c_B affect this condition only by altering the values of W^A, V_n^A, V_0^A, and q_{k^*-1}. Let $\Delta \equiv p - p_n - (c_A + c_B)$, and note that Δ increases in $p - p_n$ and decreases in $c_A + c_B$. We show that as Δ increases, this condition becomes harder to satisfy, so that A should shift to attacking sooner, decreasing equilibrium k. Re-arranging terms and using our assumption that δ is close to 1, we arrive at:

$$\rho_{k^*-1}(1-\epsilon)\sigma\left(V_0^A - W^A\right) \geq (1-\rho_{k^*-1})\lambda\frac{\Delta}{1-\delta} + \rho_{k^*-1}\epsilon\lambda\frac{\Delta}{1-\delta} \tag{3}$$

Plainly the right side increases in Δ, making the condition harder to satisfy. The only remaining issue is to establish how the left side varies in the parameters, in particular $V_0^A - W^A$.

First we show that if p increases or p_n decreases, then holding k^* constant, V_0^A must decline relative to W^A. By Proposition 4, each equilibrium offer is set to either the minimum offer that $B1$ will accept, or 1 if the former is not feasible. Holding the equilibrium offers constant, as p increases or p_n decreases, B's war value decreases relative to his possible pre-war continuation values (since these include the possibility of eventually getting V_n^B). This means the equilibrium offers will also shift upward (if they were not already 1) since $B1$ would accept less generous offers than before. However, the new offers cannot bring A back to her original value of $V_0^A - W^A$. The change in these offers does not change the value of the game, so that the new offers at most simply shift back to A the share of the surplus she was previously enjoying, given that she is facing $B1$. If any of the budget constraints (i.e., that the offers cannot be more than 1) bind, then A will get less of the surplus than before, so that $V_0^A - W^A$ goes down even with the stingier offers to B. Moreover, there is a positive probability that at some future point A will be (unknowingly) facing $B2$ rather than $B1$, and there is no way for A to make offers stingy enough to compensate for $B2$'s greater probability of eventually getting V_n^B without causing $B1$ to reject the offers. Thus $V_0^A - W^A$ must decline as p goes up or p_n goes down, and the left side of condition 3 declines, making it harder to satisfy.

Now consider a decrease in c_A. This has no effect on W^B and V_n^B, so that the equilibrium offers to B will not change. It also does not alter V_n^A, but W^A increases. Since V_0^A is an expectation taken over discounted sums of the equilibrium offers and the eventual outcome of either V_n^A or W^A, then V_0^A cannot rise any faster in c_A than W^A does. Thus, $V_0^A - W^A$ either stays the same or declines as c_A rises, making condition 3 (weakly) harder to satisfy.

Finally, consider a decrease in c_B. Observe that, holding the equilibrium offers constant, B's future continuation values will decline relative to W^B. Both W^B and V_n^B increase as c_B decreases, but by the same amount, so that there is no rise in V_n^B relative to W^B, but the value B will receive from offers prior to war or proliferation declines relative to W^B. Thus, to satisfy $B1$, A must make more generous offers. W^A does not depend on c_B, but V_A^n falls as c_B decreases. Since V_n^A falls and the new equilibrium offers must be more generous to B, V_0^A must decline as c_B decreases, making condition 3 harder to satisfy. □

Proposition 6

First consider a small-enough increase in λ. This will leave k^* unchanged, but directly affect the probabilities of alternative paths through the equilibrium. In every period prior to proliferation or war, the probability that B's investment will successfully produce nuclear weapons in the next period is either $\epsilon\lambda$ or λ, so that both these go up with an increase in λ. War occurs only if B is caught in the second stage or if A reaches k^* consecutive signals and attacks, each of which is less likely since B is more likely to have produced nuclear weapons before either occurs, and no more likely to have reached the second stage when A gets a stage signal. Mistaken war is equally likely, since it occurs only if B remains in the first stage after A has received k^* consecutive null stage signals, the probability of which is not affected by λ (only by ϵ and σ). The expected length of the game decreases since the expected time B's program will remain in the first stage is unchanged, while the expected time it will remain in the second stage decreases, and the probability of transitioning between the two is unaffected.

In simulations, there are cases in which a large-enough increase in λ increases k^*. This occurs only when ϵ is relatively low and close to λ. From condition 2 above, we can see why this would be true. Increasing λ has three effects on the condition. First, it shifts probability weight on the right side of the condition from W^A to V_n^A, decreasing the right side and leading A to be willing to attack sooner. Second, however, by Proposition 3, increasing λ also increases ρ_{k^*}, so that probability weight is shifted from the first bracketed term to the second. Since $V_n^A < W^A \leq V_0^A$, the first bracketed term is lower than the second, so that the reduction in ρ_{k^*} increases the right side of the condition. Third, increasing λ, by increasing the probability the game ends in (costless) proliferation rather than costly war, increases the surplus, and thus can increase V_0^A, as A is able to make less generous offers to B since B gets more out of investment. The last two effects will only overwhelm the first if ϵ is sufficiently low, so that the first bracketed term is lower than the second by enough, and sufficiently close to λ, so that the increase in ρ_{k^*} is high enough. Figure A1 demonstrates this possibility for a particular case. Otherwise, large-enough increases in λ decrease k^*.

The indirect effect of a large-enough increase in λ through k outweighs its direct effects on the probabilities of the various outcomes when the two effects are in opposite directions. To see why, first observe that as k^* decreases, the incremental decrease in the overall probability of proliferation increases: the fewer chances B has to invest before A attacks, the bigger the effect of reducing the number of chances. This implies that each jump down in the probability of proliferation (and jump up in that of war) is larger as k^* is further reduced. Moreover, the condition for A to prefer attacking even sooner (condition 3) becomes tighter as k^* goes down. This occurs because, as k^* decreases, the probability of war goes up and the surplus of the game decreases, so that V_0^A gets smaller and q_{k^*} declines. Thus, A will attack sooner in reaction to a smaller and smaller increase in the probability that B will eventually proliferate. □

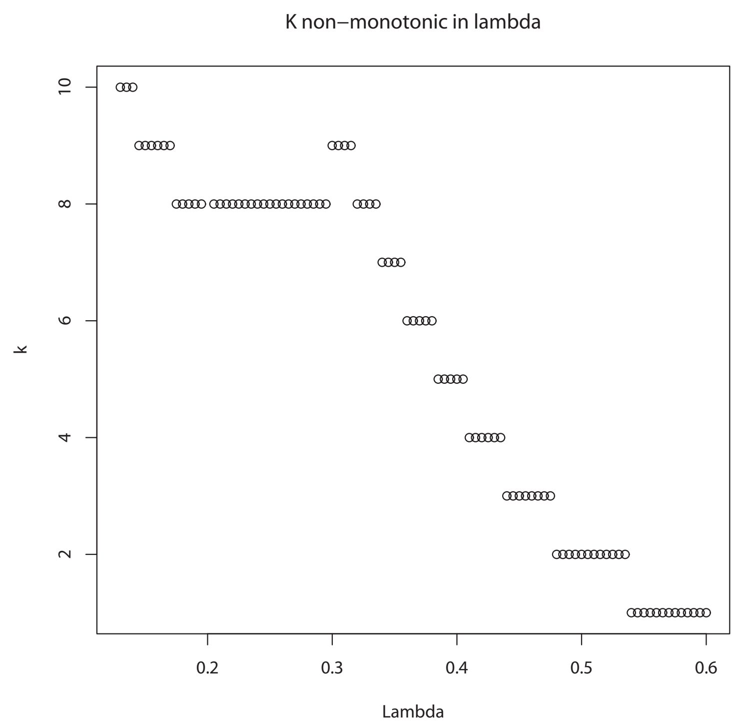

Figure A1: Equilibrium k increases in λ ($\epsilon = .2$, $\sigma = .825$, $\delta = .8$, $p = .9$, $p_n = .55$, $c_A = .05$, $c_B = .01$).

Proposition 7

First consider a small-enough increase in ϵ. This will leave the equilibrium value of k unchanged, but directly affect the probabilities of alternative paths through the equilibrium. In any period in which B's program has reached the second stage, the value of ϵ has no effect on the probabilities of various outcomes. In any period in which B's program is in the first stage, increasing ϵ will increase the probability of immediate proliferation and of moving to the second stage. Observe that, in any given period prior to proliferation or war, the probability of the game subsequently ending with proliferation is higher if B's program is second-stage than if it is first-stage. This is implied by the proof of Proposition 2: the higher chance of proliferation upon reaching the second stage is what motivates B to always invest while his program remains in the first stage. Thus, increasing ϵ must increase the probability of proliferation overall, and because the game can only end in proliferation or war, the overall probability of war must decrease. The probability of mistaken war decreases, because this can only occur if B's program has failed to advance to the second stage by the time A has received k^* consecutive null stage signals, and raising ϵ makes it more likely B will have reached the second stage by this point, and less likely that this point will be reached at all. Finally, since the expected length of the game is obviously shorter if B's program is in the second stage than if it is in the first, increasing ϵ also shortens the game in expectation.

Large-enough changes in ϵ will affect k^*. From condition 2 above, increasing ϵ lowers the second bracketed term, and transfers probability weight through its effect on ρ_{k^*} to the first bracketed term, which is less than the second, further reducing the right side. However, by increasing the probability the game ends in (costless) proliferation rather than costly war, increasing ϵ can result in increased V_0^A, as A is able to make less generous offers to B since $B1$ gets more out of investment. This latter effect is modest, since the requirement to satisfy $B1$ with these offers restricts B's ability to fully compensate A for the increasing probability of proliferation (A is fully compensated if she faces $B1$, but she is more likely to be facing

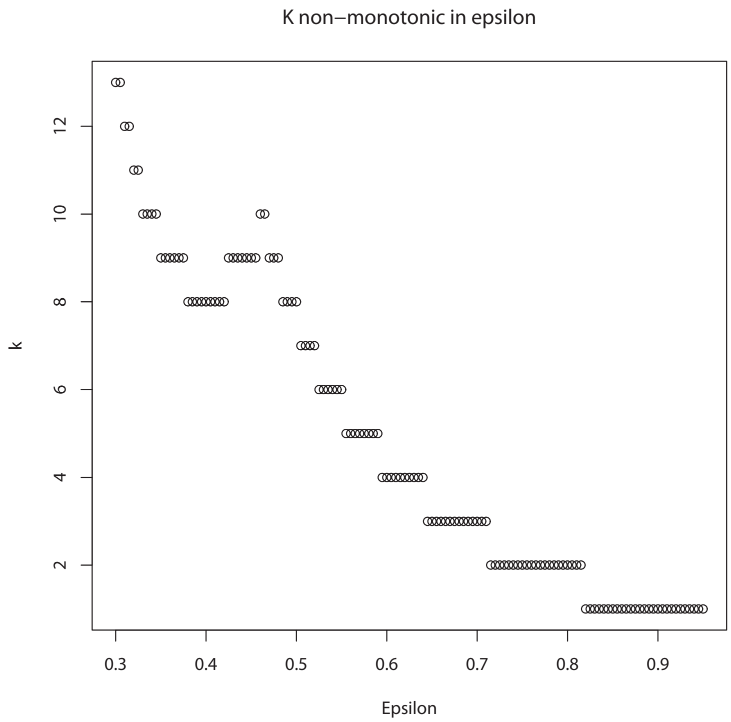

Figure A2: Equilibrium k is non-monotonic in ϵ ($\sigma = .3$, $\lambda = .3625$, $\delta = .6$, $p = .9$, $p_N = .55$, $c_A = .05$, $c_B = .01$).

$B2$ in a given period). Moreover, B's ability to compensate A in this way is limited by the constraint that $q \leq 1$. Thus, while there are cases in which an increase in ϵ will increase k^*, these are very rare in simulations. Figure A2 offers an example of such a case.

Otherwise, large-enough increases in ϵ decrease k^*, and their indirect effect through k outweighs their direct effects on the probabilities of the various outcomes, by the same argument as in the proof of Proposition 6. □

Proposition 8

First consider the direct effects of increasing σ, holding k^* constant. Proliferation directly from the first stage becomes more likely, as there is a greater chance of A receiving a stage signal while B's program is in the first stage, resetting A's estimate and giving B more chances to invest, which also increases the expected length of the game. However, proliferation from the second stage becomes less likely, as $B2$ is more likely to get caught and attacked before his program succeeds, which also shortens the expected length of the game. Which of these effects dominates depends on σ and k^*: the lower these are, the higher the proportion of the game's length B's program will be in the first stage and the more likely mistaken wars are, and the more increasing σ will tend to directly increase the probability of proliferation and the expected length. To demonstrate this possibility, consider the case when $k^* = 1$. Then the probability of proliferation is $P = \epsilon\lambda + (1-\epsilon)\sigma P = \frac{\epsilon\lambda}{1-(1-\epsilon)\sigma}$, which increases in σ. As k^* rises, B's program is more likely to reach the second stage before A attacks, and so the decrease in the probability of proliferation from the second stage and in the expected length comes to dominate. To demonstrate this, consider the case when $k^* = \infty$ (i.e., where A never attacks based only on suspicion, but only once a second-stage signal is received). It is easily shown that the probability of proliferation is $P = \frac{\lambda}{\lambda+\sigma-\lambda\sigma}$, which decreases in σ. However, the direct effect of increasing σ on the probability of mistaken war is always negative, since mistaken war can only happen if A receives k^* consecutive null stage signals, which is obviously less likely as the probability of a signal goes up.

To see the effect of σ on k^*, consider condition 2. Increasing σ shifts probability weight from W^A to V_0^A, increasing the right side. However, it can also reduce V_0^A by forcing A to make more generous offers to satisfy $B1$, since B is less likely to eventually proliferate. In some cases, the latter effect can dominate, so that increasing σ decreases k^*, but it is extremely rare in simulations. Figure A3 offers an example of such a case.

Otherwise, large-enough increases in σ increase k^*. This decreases the probability of mistaken war by Lemma 2, in agreement with σ's direct effect. In simulations, the direct effect of σ generally dominates when it disagrees with the indirect effect through k^*, so that larger increases in σ reduce the probability of proliferation and expected length of the game.

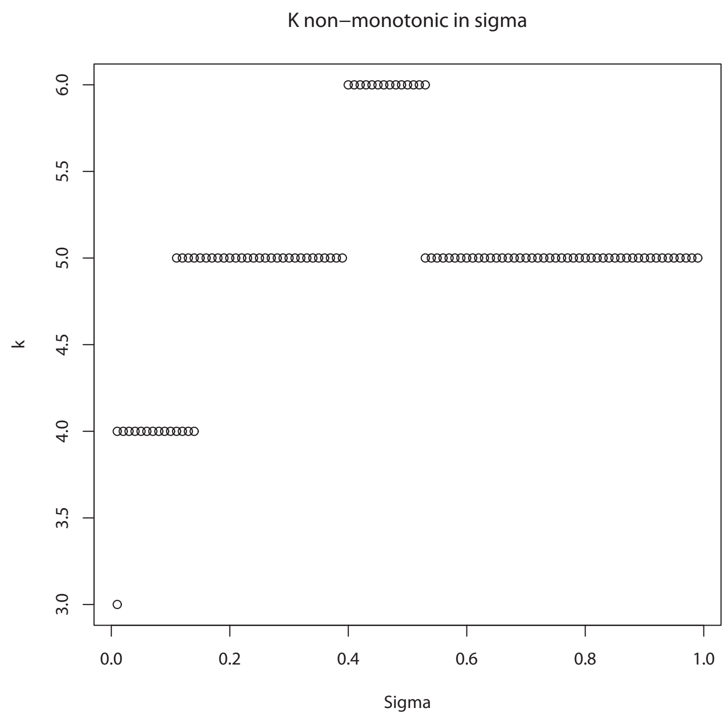

Figure A3: Equilibrium k is non-monotonic in σ ($\epsilon = .4$, $\lambda = .0625$, $\delta = .9$, $p = .9$, $p_N = .55$, $c_A = .05$, $c_B = .01$).

References

Albright, David and Paul Brannan. 2008. *The Al Kibar Reactor: Extraordinary Camouflage, Troubling Implications*. Institute for Science and International Security.

Braut-Hegghammer, Målfrid. 2011. "Revisiting Osirak: Preventive Attacks and Nuclear Proliferation Risks." *International Security* 36(1):101–132.

Burr, William. 2009. "U.S.-Iran Nuclear Negotiations in 1970s Featured Shah's Nationalism and U.S. Weapons Worries." National Security Archive Electronic Briefing Book No. 268.

Central Intelligence Agency. 1963. "The Advanced Weapons Programs of the UAR and Israel." SNIE 30-2-63.

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

Devroy, Ann and Barton Gellman. 1993. "U.S. Attacks Industrial Site Near Baghdad: MiG Downed As Gulf Allies Display Might." *The Washington Post* .

Fuhrmann, Matthew and Sarah E. Kreps. 2010. "Targeting Nuclear Programs in War and Peace: A Quantitative Empirical Analysis, 1941–2000." *Journal of Conflict Resolution* 54(6):831–859.

Grunden, Walter E. 1998. "Hungnam and the Japanese Atomic Bomb: Recent Historiography of a Postwar Myth." *Intelligence and National Security* 13(2):32–60.

Koch, Andrew and Jeanette Wolf. 1997. "Iran's Nuclear Procurement Program: How Close to the Bomb?" *The Nonproliferation Review* 5(1):123–135.

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

National Intelligence Council. 2007. "Iran: Nuclear Intentions and Capabilities." http://graphics8.nytimes.com/packages/pdf/international/20071203_release.pdf.

Nuclear Threat Initiative. 2011. "Iran Nuclear Chronology." http://www.nti.org/media/pdfs/iran_nuclear.pdf?_=1316542527.

Pollack, Kenneth M. 2002. *The Threatening Storm: The Case for Invading Iraq*. Random House Inc.

Sadot, Uri. 2015. "Osirak and the Counter-Proliferation Puzzle." Working Paper.

Thomas, Gordon. 2015. *Gideon's Spies: The Secret History of the Mossad*. 7th ed. St. Martin's Press.

Walsh, James Joseph. 2001. Bombs Unbuilt: Power, Ideas, and Institutions in International Politics PhD thesis Massachusetts Institute of Technology.

Walsh, Jim. 1997. "Surprise Down Under: The Secret History of Australia's Nuclear Ambitions." *The Nonproliferation Review* 5(1):1–20.