Journal of Conflict Resolution
56(4) 651-674

DOI: 10.1177/0022002712445740
http://jcr.sagepub.com

Arms Diffusion and War

Muhammet A. Bas[1] and Andrew J. Coe[1]

Abstract

The authors present a model of the relationship between the spread of new military technologies and the occurrence of war. A new technology could shift the balance of power, causing anticipatory war as one side tries to prevent the other from obtaining it. When one side already has it, war is more likely when the shift in power is large, likely, and durable. When neither side has it, war is more likely when the expected shift is asymmetric (e.g., one side is more likely to get it) and when the two sides fear that a war will occur once one of them has it. The authors illustrate the model with historical examples from the spread of firearms (the Musket Wars in precolonial New Zealand) and of nuclear weapons (the end of US nuclear monopoly and the 1967 Six-Day War). A broader implication is that major power competition can unintentionally cause wars elsewhere.

Keywords

preventive war, proliferation, bargaining models, nuclear weapons

Sitting at his desk in June 1946, President Harry S. Truman wrote a note to himself: "Get plenty of Atomic Bombs on hand—drop one on Stalin, put the United Nations to work, and eventually set up a free world."[1] Though the United States was then the sole possessor of nuclear weapons, Truman was well aware that the Soviets were hard at work trying to copy the US invention and that other states might follow suit. Indeed, only three years later, the Soviets would test their

[1]Department of Government, Harvard University, Cambridge, MA, USA

Corresponding Author:
Muhammet A. Bas, Department of Government, Harvard University, Cambridge, MA, USA
Email: mbas@gov.harvard.edu

own atomic bomb, to Truman's great dismay. Yet, Truman did nothing to stop it. Why?

The invention of military technology is often followed by its diffusion around the world, and the introduction of new weapons into new locales is sometimes attended by war. The spread of firearms to precolonial populations was concomitant with intense warfare among these groups, and the spread of nuclear weapons has also been marred by occasional wars. Were some of these wars *caused* by arms diffusion? If so, why did wars result in some cases but not others? How does the occurrence of war depend on the characteristics of the new weapons, and of the states to which they could spread? In particular, why was Truman content to bide his time while the Soviets got the bomb?

We explore these questions using a game-theoretic model of bargaining between two states in an environment in which new technologies are spreading. In the model, the acquisition of a new technology can shift the balance of military power between the two states. The likelihood of a state acquiring the new technology depends on its technological sophistication and its competence in espionage; the magnitude of the resulting shift depends on the specific technology and on the preexisting military capabilities of the two states. Over time, a state may become more likely to acquire a particular technology, as the weapons' external spread, the state's spying efforts, and/or its internal development progress. And the introduction of new technologies may recur over time.

We find that states' inability to commit to not taking advantage of their new weapons can, under certain conditions, lead to a preventive war aimed at stopping a state from getting the new technology. States' expectations about the future are critical in determining whether war occurs. We show that war is more likely when the expected shift in relative power is large, likely, and durable, in line with the arguments of scholars of the preventive motivation for war (Taylor 1954; Levy 1987; Copeland 2000). More surprisingly, we find that war can occur when neither, one, or both have acquired a new technology: the commonsense intuition that the preventive motivation for war will be highest when one state, but not the other, possesses the new weapons is true only under certain conditions. We also demonstrate that preventive war is more likely when states anticipate a greater probability of war in the future. Scholars of preventive war have long argued for this proposition, but it has not appeared previously in the bargaining literature on war.

We illustrate the model by applying it to several examples drawn from the history of the spread of firearms and of nuclear weapons. Each of these technologies have been highly consequential for states' military capabilities, and so each is a prime candidate for generating preventive wars. We first look at the introduction of firearms to Maori tribes in New Zealand in the early nineteenth century, which was followed by the aptly named "Musket Wars." This series of unusually intense conflicts was initially characterized by tribes with newly acquired muskets attempting to destroy or enslave tribes without muskets—for preventive reasons, we argue. Next, we examine the end of US nuclear monopoly, with the spread of nuclear weapons to the Soviet Union in the late 1940s. Our model suggests, and there is some evidence to support, that US hesitation to forcefully protect the monopoly was driven in part by the (violated) expectation that it would be a long time before the Soviets could successfully copy the bomb, in contrast to other accounts of this period that emphasize US lack of preparedness for or moral qualms about war. Finally, we turn to the 1967 Six-Day War, between Israel and Egypt (among others). Our model predicts that preventive motivations with respect to Israel's nuclear program should have played a part in the origins of the war, supporting a recent and controversial line of historical work on this case.

In related studies, Fearon (1995) and Powell (1999) also analyze exogenous shifts in the balance of power, but their models treat these shifts as deterministic and perfectly anticipated, whereas our model allows them to be stochastic and hence uncertain. This gives rise to different empirical implications: states' war decisions turn on *expectations* about future shifts rather than the shifts that will actually occur, and so states may sometimes consider or launch preventive attacks even when no shift later occurs, and they may forego prevention even when a large shift later occurs.

Baliga and Sjöström (2008) and Debs and Monteiro (2010) deal explicitly with the situation in which one state can invest in acquiring a new military technology, but their models allow only one opportunity for the balance of power to shift, whereas our model allows both multiple opportunities for a given shift and multiple shifts as first one side, then the other tries to acquire a new technology. The cost of this additional generality is that, for tractability, we assume the shifts are exogenous, and as a result cannot assess the role of secrecy and ambiguity in a state's decision to *pursue* the new technology, which is the central concern of the cited models. The benefit is that we enable study of the role of secrecy about the *progress* a state is making toward acquiring a new technology, given that its decision to do so is known, shedding light on an aspect of secrecy ignored by earlier models. Additionally, we can analyze the situation in which neither side has the technology and also study the effects of recurring technological diffusion. We thus view our work as complementary to these efforts.

Finally, Horowitz (2010) and Rosen (1991) study the determinants of successful innovation, incorporation, and application of new technology by nations' militaries. This research focuses on the *causes* of diffusion (as well as innovation), whereas our study is focused on its *effects*. Our model simply parameterizes the ability of states to copy a new technology, so that we may analyze the effect of this and other parameters on the potential for preventive war.

In the next section, we describe the model and our assumptions. The subsequent section characterizes the equilibria of the model and derives comparative statics. We then turn to the historical illustrations of the model. The final section offers some broader implications of our analysis.

Setup of the Model

We model the interaction between two states (or groups), A and B, as they bargain repeatedly over the division of a contested stake, represented by the unit interval.

The players alternate between making and responding to proposals for revision of the status quo over the course of potentially infinitely many periods, with one offer per period. If the receiving player accepts the proposal, the revision is implemented immediately. If he rejects the offer peacefully, the status quo division is maintained. Finally, the receiving player can respond by starting a war, in which case the game ends with a costly lottery. The value of this lottery to each player depends on the balance of military power between them, represented by A's probability of victory in war. The winner receives the entire stake in this and all future rounds; the loser gets nothing. Regardless of who wins, each player pays a positive cost, c_A and c_B respectively, in this and all future periods, which represents the destructiveness of the war.[2]

In any given round, a peaceful response to an offer (whether acceptance or rejection) is followed by a move by Nature which determines whether a new military technology diffuses to each player. In the Maori case, this can be thought of as the arrival (or not) of a European ship bearing firearms for trade; in the case of the initial spread of nuclear weapons, this is the success (or not) of the Soviet Union's development program. We assume that this process is exogenous—the players have no control over the occurrence of diffusion.

This assumption is reasonable in the empirical cases discussed here, and in many others, because two conditions are met. First, the cost in terms of resources of acquiring the new technology must be outweighed by its benefit in terms of increased bargaining power over the contested stake, so that a state would try to acquire the technology, given the freedom to do so. The Maori tribes and the Soviet Union and Israel had much to gain over their adversaries from acquiring firearms and nuclear weapons (respectively). And these were not unaffordable: the Maori had goods the European traders desired in abundance, and the Soviet Union's size and Israel's wealth meant that each could easily devote the necessary resources to developing nuclear weapons. Second, arms control agreements must be unenforceable, usually because one side cannot monitor the other's acquisition efforts well enough to assure compliance, so that states could not be persuaded to eschew the technology voluntarily. The Maori could not predict when and where European traders might next appear, and so had no ability to enforce agreements to limit their purchases of firearms when an opportunity arose. And the US and Egyptian capabilities for monitoring the Soviet and Israeli nuclear programs (respectively) were quite rudimentary, as we will document in the next section. These two conditions ensure that the players would pursue the new technology given the chance, so that we can take diffusion as exogenous.[3]

If a player receives the technology, the balance of power shifts. Before either has the technology, the balance of power (A's probability of victory in war) is p_{00}; when A has it but B does not, the balance is p_{10} and similarly for p_{01}; and when both have it, the balance is p_{11}. Naturally, we assume that $p_{01} < p_{00} < p_{10}$ and $p_{01} < p_{11} < p_{10}$, so that unilateral possession of the technology is advantageous in war, relative to mutual acquisition or mutual lack of the technology. However, we allow p_{00} to differ from p_{11}. If, for example, state A has conventional forces superior to those of state B, but both have acquired nuclear weapons, then conventional arms may be of reduced importance in determining the overall balance.

We assume that the balance of power (variously p_{00}, p_{10}, p_{01}, and p_{11}) is also exogenous: the players cannot affect it by arming or transferring strategic resources between them. It has been shown elsewhere that each of these measures for affecting the balance of power can independently cause war (see Powell 2006)—we rule them out to isolate the role of diffusion in war and simplify the analysis. However, their inclusion would not qualitatively alter our results. We show that diffusion can only cause war if the shift in the balance of power resulting from diffusion is large enough, and will not cause war if the shift is too small. Introducing the possibility of arming or strategic transfers might lead to wars even when the shift from diffusion is small, but these wars would not be caused solely by diffusion. By contrast, arming or transfers might enable the actors to compensate for diffusion—for example, one side could arm more heavily in response to the other's acquisition of a new technology, lessening the net change in power. But arming or transfers suffice to avoid war due to diffusion only when the (qualitatively) new technology is relatively easy to match with quantitative improvements in arms or strategic resources. We can safely exclude these measures from our model because this condition did not hold in the cases we consider: a Maori tribe could not easily match an opposing tribe's acquisition of firearms with more clubs (or firearms, if it already possessed them), and the United States and Egypt could not easily compensate for their opponents' acquisition of nuclear weapons.

The probability of diffusion depends on whether the other state has acquired the technology. When neither has it, A's probability of receiving it in the next round is λ_A, and similarly for λ_B. These are allowed to differ as a proxy for different levels of susceptibility to diffusion between the two states. For example, two states with economies of differing technological sophistication may have different probabilities of initial diffusion. If A already has the technology, then the probability that B will receive the technology in the next round is λ_{AB}, and similarly for λ_{BA}. We assume $\lambda_{AB} \geq \lambda_B$ and $\lambda_{BA} \geq \lambda_A$: once one state has the technology, the other state is (weakly) more likely to receive it. This is plausible because the latter state now knows that technology is available and feasible and this may guide its efforts to duplicate it. Moreover, once one state has the technology, the other may employ espionage against the former and eventually steal the technology.

As a simplification, we typically assume that the various λs do not change over time. That is, if a player does not receive the new technology in this round, then his chance of acquiring it in the next is the same, so long as the other player does not receive it in this round. So, the per-period chance of diffusion is not cumulative, even though the ex ante probability of diffusion occurring over the next several periods in increasing in the number of periods considered. This assumption is less applicable in the case of the spread of nuclear weapons. We relax it in Cumulative Diffusion' subsection, allowing λ to increase over time, and discuss the differences this introduces.

Finally, we assume that there is only one new technology available. That is, once a player has received the technology, there is no *other* new technology to be acquired that could again shift the balance in his favor. This assumption greatly simplifies the exposition of the analysis, but it is unrealistic. We show in the supplemental appendix that relaxing this assumption does not alter the qualitative character of our results.

A player's utility is assumed to be linear in present consumption, defined as the player's share of the contested stake. Settlements are labeled by the share going to A: the settlement in which A receives q and B receives $1 - q$ is called q. The players are assumed to discount future consumption at a common rate of $\delta < 1$ per period. Players' preferences and all the exogenous parameters of the game are common knowledge.

Equilibrium Analysis

We use backward induction to find the subgame perfect equilibria (SPE) of the model. We start with the subgame in which both players have the technology, termed "mutual armament," then turn to the prior subgame in which one player has the technology but the other does not, called "diffusion."[4] Next, we relax the assumption that the probabilities of diffusion are constant, analyzing the case in which they increase over time. Finally, we examine the subgame in which neither player has the technology, labeled "introduction." Proofs of all the propositions appear in a supplemental appendix, available at the journal's website.

Mutual Armament Subgame

We begin with the subgame where both players have the technology. Proposition 1 characterizes the unique outcome of any SPE of this subgame. This result is similar to that of standard models of shifting power, so we only remind readers of the intuition (Fearon 1995; Powell 1999).

Proposition 1: Suppose that both players have the technology and that the status quo division is q. No war will occur. If $q < p_{11} - c_A$, then a permanent revision to $q' = p_{11} - c_A$ will be immediately agreed. If $q > p_{11} + c_B$, then a permanent revision to $q' = p_{11} + c_B$ will be immediately agreed. If $q \in [p_{11} - c_A, p_{11} + c_B]$, there will be no revisions.

Intuitively, once both players have the technology, there is no longer any uncertainty or any possibility of a future shift in the balance of power. The players can thus find a settlement that both prefer to war, and peace will prevail. Furthermore, aside perhaps from a single, immediate revision to bring the status quo division of the disputed stake into line with the balance of power, there is no cause for further revisions. The key comparative static is that the higher p_{11} is, the better A (and the worse B) will expect to do once both players have the technology.

Diffusion Subgame

Backing up, we assume that A has the technology, but B does not. Proposition 2 characterizes the unique outcome of any SPE of this subgame.[5] As we explain, the mechanism that leads to war here is the same as in the standard models of shifting power, but the interpretation is different.

> *Proposition 2:* Suppose player A has the technology, but B does not. If $p_{10} - c_A + \frac{\delta}{1-\delta}\lambda_{AB}(p_{10} - p_{11}) > 1 + \frac{\delta}{1-\delta}\lambda_{AB}(c_A + c_B)$, then there is war. Otherwise, there is peace, and at most one revision to the status quo, which is immediately agreed.

The root cause of war is that B faces a commitment problem. He is likely to receive the technology at some point in the future, and when he does, the balance of power will shift in his favor. By Proposition 1, he will then demand and receive a revision that will bring the status quo into line with the new balance of power, leaving him better off and A worse off. B cannot commit himself ahead of time to not taking advantage of the technology once he acquires it, so he must try to buy A off by making a concession now in order to compensate him for the adverse shift to come; otherwise, A would start a war in order to prevent the shift. If the shift is sufficiently likely and sufficiently large, then the largest concession B could make—offering the whole stake to A until the shift occurs—is not enough to placate A, and war will occur. However, if the shift is not so likely or so large, then B can successfully compensate A and the game is peaceful.

War can occur in this subgame because A has reason to fear the future: B will eventually receive the new technology, and A's temporary advantage deriving from its sole possession of the technology will be lost. Whether the loss of this advantage will compel a preventive war depends first on how large the anticipated loss will be. Proposition 2 implies that war can only occur when the anticipated shift in power (i.e., $p_{10} - p_{11}$), is larger than the total cost of war ($c_A + c_B$). If the shift is not this large, then war is so destructive that its costs outweigh the benefits to A of locking in his temporary advantage.

However, unlike in the standard models in Fearon (1995) and Powell (1999), even a very large shift in the balance of power due to the diffusion of the new technology may not cause war. If the shift is larger than the costs, whether war occurs or not also depends on how likely the shift is to occur. It is clear from the condition in Proposition 2 that an increase in the likelihood of the shift (λ_{AB}) makes war more likely. A has more reason to launch a preventive war to the extent that his advantage will be lost sooner. But conversely, if the chances of B acquiring the new technology are very low, then A will not start a war even if diffusion would be highly consequential for the balance of power, because diffusion will not happen any time soon. We delve further into this implication by examining what happens if the chance of diffusion increases over time.

Cumulative Diffusion. The analysis in the previous section assumed that the per-period probability of diffusion (e.g., λ_{AB}) is constant over time. In reality, diffusion is often increasingly likely to occur over time. For example, a state's efforts to acquire nuclear weapons usually have an increasing chance of success as time goes by and progress in indigenous development, external assistance, and/or espionage occurs. In this section, we study a variant of the model in which the probability of diffusion increases over time, and show that the corresponding incentives for preventive war also grow over time. We show that a state facing the possibility of diffusion to another might choose to delay a preventive attack, putting off the costs of war until the other side is close to acquiring the technology, even though this poses some risk of diffusion occurring unexpectedly early. The analysis also reveals why the acquiring state might conceal its progress in duplicating the technology.

Suppose as in the previous section that player A has the new technology, but B does not. The game here is identical to that of the previous section, except that λ_{AB} is now an increasing function of time. To economize on notation, suppose that the game starts at time $t = 0$ and that at time t, $\lambda_{AB} = \lambda_t$. We assume that λ_t is strictly increasing for $t \in \{0, \ldots, n\}$ for some $n < \infty$, and is constant at λ_n thereafter.[6]

> *Proposition 3:* Suppose player A has the technology, but B does not, and $p_{10} - p_{11} > c_A + c_B$. So long as B does not get it in the meantime, the value of war relative to peace for A increases over time until period n, and is constant thereafter.

As in the previous section, the requirement that $p_{10} - p_{11} > c_A + c_B$ ensures that there is at least some incentive for preventive attack: it means that the advantage A has due to sole possession of the technology is large enough to make it worth contemplating a costly war to preserve it. If this condition did not hold, A would not attack regardless of the value of λ_t—even if it were 1, so that B was *certain* to get the technology if A didn't act to stop it. Given the condition, the proposition shows that the case for A to attack gets stronger as time passes and λ_t increases.

Figure 1 illustrates the intuition for the result. The values of λ_t control when B is expected to get the new technology. If the first few values of λ_t are low enough, A will expect that it will be a long time before B ends its monopoly, and so will have little reason to worry about it. In particular, A will not attack in the early periods because he would rather put off the costs of preventive war until a later date. However, there is a trade-off between avoiding the costs of war in the current period and taking the risk that B gets the technology in the meantime. As time passes and λ_t grows, it becomes more likely that B will get the technology very soon, and A has more cause to worry about the imminent loss of its advantage. If λ_t gets high enough, those worries will overcome A's desire to postpone the costs of war and A will attack.

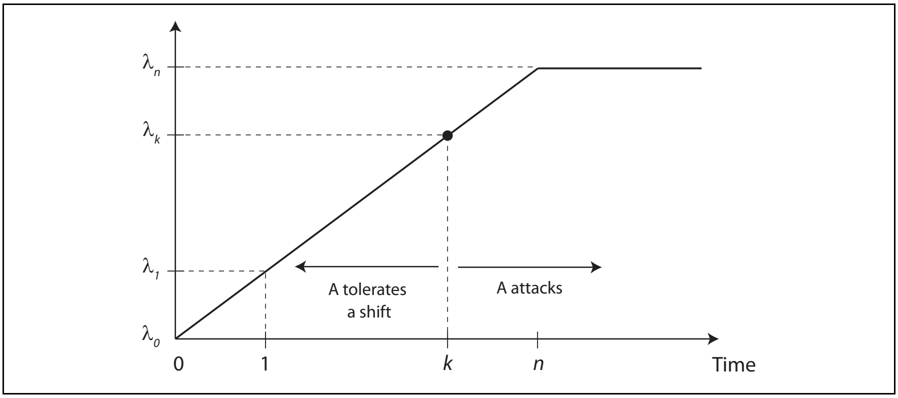

Figure 1. Gradually increasing λ_t and delayed attack

An implication of this is that *B* might get "lucky." That is, it might be that *B* acquires the technology faster than expected, while λ_t is relatively low, and before *A* is stirred to act. This explains why states are sometimes surprised by diffusion, and why the state that already had the technology would "let it happen." In empirical cases where diffusion occurs without preventive war, the explanation could be that the costs of preventive war were too high, so that *A* would not attack even if diffusion was imminent and would let diffusion occur. But our analysis suggests an alternative explanation: that *A* would have attacked but missed its chance due to *B*'s luck, and so had no choice but to tolerate the shift. We develop this further in our discussion of the end of US nuclear monopoly.

The analysis also shows why a state trying to acquire the technology might conceal its progress. Especially for difficult technologies, concealment generates uncertainty on the part of the state that has it about when the other will get it and can thus lower incentives for prevention. To see why, consider two extreme cases. In the first, *A* can perfectly observe *B*'s progress. In the second, *A* sees nothing until the shift occurs.

In the first case, since *A* has the technology and knows the steps involved in its development, he will be able to tell exactly when *B*'s acquisition is imminent. That is, he can watch as *B*'s efforts progress and see when it reaches the verge of success. At that point, the probability of diffusion is 1, and so *A* has the strongest possible incentive to attack then, because if he does not, the balance of power will shift with certainty. Obviously this does not bode well for *B*. If *A* chose to attack, then *B* has no chance of acquiring the technology and enjoying the benefits of the resulting shift in power, and B will also have to pay the cost of war.

If instead *A* could not observe anything, he would always be uncertain about how far *B* was from getting the technology. He would have to use his own experience with the technology and his knowledge of *B*'s technological sophistication to estimate the probability that B would acquire the technology in the next period, for each period: that is, he would have to estimate the λ_ts. A more difficult technology would have more steps involved in its acquisition, more possibilities for wrong turns, and more opportunities for delay as various elements were mastered. It would thus lead A to assign a more diffuse distribution over when B might be successful. Individual values of λ_t would be lower and would rise more slowly in t. Certainly, λ_t would not equal 1 in any period. Thus, at any time, the incentives for preventive attack would be lower than in the previous case at the (known by A) time at which diffusion was imminent.

It might be the case that the distribution was so spread out that A would not ever feel an urgent need to attack, no matter how long it took B to be successful. Or, consistent with proposition 3, there might come a time when A would think that diffusion was close enough (i.e., λ_t was high enough) to warrant preventive attack. But the time that would pass before that point would constitute time for B to get lucky. Thus, relative to the first case, B would have a higher chance of acquiring the technology and gaining the corresponding advantage, as a well as a lower chance of having to pay the cost of war.

Most empirical cases would fall somewhere between these two extremes. A would have *some* ability to observe B's progress, and B would have *some* ability to conceal its efforts. But the comparison of the two extremes makes clear why secrecy would be beneficial for B (and harmful to A) and surveillance beneficial for A (and harmful to B).

Previous models of shifting power did not produce these results because they assume either that the shift is deterministic or that it can only occur at one point in time (Fearon 1995; Powell 1999; Baliga and Sjöström 2008; Debs and Monteiro 2010). When this is true, B has incentives to try to conceal the *existence* of its efforts at acquiring the new technology from A, but there is no *progress* to speak of, so that once A knows of B's program, there is no more reason for secrecy. Either A attacks or he does not—there is no possibility of watching and waiting to attack until B is close to success, no possibility of B getting lucky in the meantime, and no possibility of A regretting its delay.

Our model explains why one state might simply watch and wait as another attempts to acquire a consequential new technology and to conceal its progress in doing so. The costs of preventive war are not worth suffering until diffusion is near. Because of this, the acquiring state will try to conceal its progress, so that the watching state is never certain that diffusion is imminent, and might get lucky, successfully developing the technology before the watching state acts to stop it.

Introduction Subgame

Finally, we turn to the subgame in which neither player has received the new technology. We first assume that the effects of the technology and the probabilities of it diffusing to either or both players are fixed and common knowledge. We show that the outcome turns on the expected shift in the balance of power and the probability of a preventive war occurring in the future. If these are large enough, then the player who expects to lose in the future may fight now to avoid the (expected) shift and risk of future war. We then discuss what happens when this assumption is relaxed, so that players are unable to perfectly forecast the "next" new technology. We argue that uncertainty about the technological future makes war in the introduction subgame less likely.

To state the main result, we need some additional notation. Let Δp be the expected shift in the balance of power:

$$\Delta p \equiv \lambda_A(1-\lambda_B)[p_{10}-p_{00}] + (1-\lambda_A)\lambda_B[p_{01}-p_{00}] + \lambda_A\lambda_B[p_{11}-p_{00}]. \qquad (1)$$

This expected shift takes into account that, in the next period, A could acquire the technology, B could, or both could simultaneously. It can be thought of as a measure of asymmetry in the possible futures. If $\Delta p = 0$, then the future is perfectly balanced; neither player expects an advantage. However, if $\Delta p > 0$, then the future is expected to benefit A, and vice versa. The further Δp is from zero, the more sharply the future favors one player over the other.

We will assume that $\Delta p \leq 0$, so that the balance of power is expected to move against A (perhaps weakly).[7] Let $V_{ij}^A(1)$ be the value A receives in equilibrium and W_{ij}^A be the war value A would receive in the ij subgame, where $ij \in \{01, 10, 11\}$, when the subgame is entered with $q = 1$. Then $S_{ij}^A \equiv V_{ij}^A(1) - W_{ij}^A$ is the value of the ij subgame over and above A's war value in that subgame, given that A starts in possession of the whole stake. We call this A's "surplus."

The values of S_{ij}^A are determined by Propositions 1 and 2, but we offer a few observations about these values to guide the intuition. First, all these values are nonnegative; since A can always go to war, he would never accept less than his war value in any subgame. In fact, if war occurs in the ij subgame, then $S_{ij}^A = 0$. By contrast, S_{ij}^A may be as large as $\frac{c_A+c_B}{1-\delta}$ if the ij subgame is peaceful and A captures the full surplus from avoiding war for the rest of the game. In fact, we know from Proposition 1 that S_{11}^A is equal to this upper bound, because, upon entering this subgame holding the entire stake, A will concede only enough to make B indifferent between war and peace.

Finally, let $\mathbf{S}^A = \lambda_A(1-\lambda_B)S_{10}^A + (1-\lambda_A)\lambda_B S_{01}^A + \lambda_A\lambda_B S_{11}^A$. This is the *expected* surplus that A will obtain if a shift occurs in the next round, given that a war does not occur in this round. It can be thought of as the reward to A from not going to war now. Observe that $\mathbf{S}^A$ is smaller when war is more likely to occur in a future subgame—we will use this in interpreting the following proposition:

Proposition 4: Suppose neither player has the technology and $\Delta p \leq 0$. Then war will occur if and only if $p_{00} - c_A - \frac{\delta}{1-\delta}\Delta p > 1 + \delta\mathbf{S}^A$.

The proposition shows that whether war occurs depends principally on the relationship between the expected shift in power Δp and the expected future surplus $\mathbf{S}^A$,

Table 1. Example Parameter Values and Equilibrium Outcomes in Subgames (δ=.9, $c_A = c_B = .1$)

	Parameter values								Subgame outcomes		
Case	λ_A	λ_B	λ_{AB}	λ_{BA}	p_{00}	p_{10}	p_{01}	p_{11}	Introduction	*A* has it	*B* has it
1	.1	.1	.25	.25	.5	.75	.25	.5	Peace	Peace	Peace
2	.1	.1	.2	.15	.5	.9	.25	.35	Peace	War	Peace
3	.1	.1	.2	.2	.5	.75	.1	.7	Peace	Peace	War
4	.1	.1	.2	.2	.5	.9	.1	.5	Peace	War	War
5	.1	.2	.3	.25	.6	.85	.15	.5	War	War	War
6	.2	.05	.25	.1	.3	.8	.1	.6	War	Peace	War
7	.1	.25	.3	.15	.7	.8	.2	.35	War	War	Peace
8	.01	.25	.3	.05	.6	.9	.1	.7	War	Peace	Peace

since these will dominate the inequality for plausibly high values of the discount factor δ. When Δp is further from zero, *A* expects the balance of power to shift more sharply against it. And when $\mathbf{S}^A$ is smaller, there is less and less reason to wait to fight, because the future offers *A* little more than he would get from fighting now. This is most obviously the case if one or both of the diffusion subgames (where just one player has the technology) result in war—*A* receives no surplus in such subgames. Intuitively, the darker the future appears to *A*, whether because of a worse expected shift or a higher probability of later war or both, the more likely it is that *A* would prefer to fight now.

This shows that the claim made by Levy (1987) and others, that the fear of future war can be a powerful motive for preventive attack, is right, at least in the context of arms diffusion. *A* might attack when neither state has the technology, in part because of a fear that if one state got the technology first, it would then fight a war to prevent the other from succeeding it. Prior models of preventive war did not exhibit this phenomenon because they assumed that the shifts in the balance of power would always favor a particular side (Fearon 1995; Powell 1999). In equilibrium in these models, the favored side would have no incentives to launch a war, as it would expect its bargaining position to improve. As a result, the other side has no rational reason to fear future war. Our model implies otherwise because it allows for multiple shifts as first one side, then the other, obtains the new technology: opposing shifts create opportunities for both players to face incentives for preventive attack, and hence to rationally anticipate the possibility of future war.

In fact, any combination of war or peace in the various subgames can occur. Table 1 gives examples of plausible parameter values that lead to each possible combination of outcomes. Notice that we can find sets of parameter values that lead to war in the introduction subgame *even when one or both diffusion subgames are peaceful.*

This may seem a strange finding. Intuitively, we would expect that the most dangerous times are when one potential belligerent has the technology, but the other does not. The last three rows in Table 1 show that this is not necessarily the case. In the next section, we will argue that this is not an artifact of the model, using the 1967 Six-Day War as an example in which war occurred before any of the belligerents had a (deliverable) nuclear weapon.

So far we have assumed that the players could confidently predict the effects of the "new" technology and the likelihood that either would acquire it. This might be the case when the players anticipate receiving a technology that external actors already possess. They could then form expectations about the effects this technology might have on their situation based on its consequences for the external actors. For example, we expect that Israel and Egypt, having witnessed the interaction of other nuclear-armed states, would have reasonable estimates of the effects on the balance of power between them should either get nuclear weapons.

However, when players cannot observe the next technology before acquiring it, they will likely be substantially uncertain of its effects. A variant of our model, in which the introduction of new technologies can recur over time and there is uncertainty about what technology comes next, is discussed in the supplemental appendix. Here, we sketch the intuition behind the main results.

Technologies that are highly consequential, in the sense of having large effects on the balance of power between two states, are rare. The "next" technology is much more likely to have modest effects on the balance of power, insufficient to cause war regardless of who has it and who does not. Thus, the expected value of Δp over the distribution of technologies is relatively close to zero and, by Proposition 4, war will not occur. (Note that if Δp is close enough to zero, the inequality in the proposition cannot be satisfied.) To illustrate this, suppose (counterfactually) that Israel and Egypt did not know that the "next" technology was nuclear weapons. Without any other information, they would likely assume that the effect of the next technology would be modest, since this is true for most new technologies, and hence, the potential for either to acquire a new technology would not be enough to cause war.

This refines the intuition that preventive wars launched prior to either side acquiring a new technology should be rare. When there is no good information available about the next technology, uncertainty and the rarity of highly consequential technologies mean that there is little cause for preventive war. By contrast, when one side has developed a new technology, both sides can observe it and estimate its effects, and should those effects be large, war may result as one side tries to defend its monopoly. However, it may also be the case that, even though neither side has the technology, both know what the next technology is and that it will have large effects on their balance of power. There would then be substantial incentives for war.

This argument carries implications for the war proneness of different eras. In eras in which all states (either in the world or in some isolated region) possess roughly similar military technology, so that diffusion of the last highly consequential new technology is complete and the next new technology is unknown, wars due to the introduction and diffusion of new technologies will be rare. But in eras in which the next technology is known and consequential, such as the era of the spread of firearms from Europe or the current era of nuclear proliferation, wars due to the introduction and diffusion of that technology to new regions will be more common. We will return to this implication in the final section, as it has further ramifications for understanding the global externalities imposed by competition among the more technologically advanced states of the world.

To summarize, we found that war could occur even before either side had the new technology, and even if one or both situations where one side had it but the other did not would be peaceful. The introduction of a new technology is more likely to cause war when states can confidently estimate its effects, when it is expected to benefit one side asymmetrically, and when states expect that, upon acquiring the technology first, one or both states might attack the other to preserve its monopoly.

Illustrations

Next we offer three empirical illustrations of the model, drawn from the histories of the spread of firearms and of nuclear weapons. The Musket Wars, fought in New Zealand in the early nineteenth century after the introduction of firearms there, demonstrate the most basic intuitions of the model. The end of the nuclear monopoly, when the Soviet Union acquired nuclear weapons without being preventively attacked, illustrates how a rise in the probability of acquisition over time generates incentives for secrecy and can result in peaceful diffusion if the acquiring state gets "lucky" and receives the technology earlier than expected. Finally, the 1967 Six-Day War is used to show that the preventive motivation for war can arise even when neither side has acquired the technology.

The Musket Wars

In the early nineteenth century, some of the Maori tribes populating New Zealand came into contact with European ships and traded with them to obtain firearms (muskets) for the first time.[8] Soon thereafter, the newly armed tribes began launching raids against other tribes that had not yet acquired muskets, marking the beginning of the Musket Wars, a series of bloody conflicts that reduced the native population by more than a quarter, radically shifted tribal territories, and indirectly paved the way for European colonization. Unlike the intermittent, low-intensity, highly localized raids that characterized Maori warfare prior to the introduction of firearms, the Musket Wars verged on total warfare, as tribes repeatedly launched all-out campaigns, ranging throughout New Zealand, to annihilate, enslave, or permanently expel other tribes. After almost thirty years, all the surviving tribes had acquired ample muskets, and the wars ended. Our model provides a coherent explanation for many aspects of these conflicts, and so sheds light on a debate among historians about their origins.

It is well established that a tribe's possession of muskets gave it a large advantage over those armed mainly with clubs and spears (i.e., $p_{10} - p_{11}$ was large), and that the arrival of European foodstuffs, especially potatoes, considerably lessened the cost of large campaigns like those of the Musket Wars (i.e., $c_A + c_B$ was low). But counter to the arguments of Belich (1996), King (1997), and Crosby (1999), these factors are not sufficient to explain either why the wars occurred or why they were so intense. If the advantages of muskets were so well understood that they became the most desirable import in New Zealand, then why did war occur? After all, the tribes without muskets should have been willing to make concessions to placate those with, and so avoid the costs of fighting, which would probably end in the formers' defeat. And if fighting instead occurred for the same reasons it had prior to the introduction of muskets, as Ballara (2003) argues, then why were the wars so much more intense and wide ranging than before? Certainly, the presence of muskets did not necessarily mean battle would be bloody: one of the most heavily armed tribes refrained from using its muskets in internal conflict (Belich 1996, 158, 161).

The key element missing from these explanations, as our model makes clear, is that the advantage muskets offered was *fleeting*—it would last only until a musket-armed tribe's opponents themselves got guns (i.e., λ_{AB} was high). If instead the first tribe to obtain muskets could somehow be assured that the others would never get them, it would have little need to engage in war, other than perhaps an occasional demonstration of the advantages afforded by its muskets, as the other tribes would simply pay tribute of one form or another in order to appease it, indefinitely. But this was not so, for the Maori had good reason to expect that the Europeans would exploit other opportunities for trade and that the other tribes would thereby acquire firearms, as indeed eventually happened.

Our model thus explains the character of the Musket Wars and why it differed from that of precontact warfare among the Maori. Proposition 2 implies that war occurred because muskets bestowed a large, but short-lived, advantage on their initial possessors over those lacking them, and because the arrival of potatoes lessened the cost of cementing this momentary opportunity. Exploiting the advantage to the fullest meant attacking any and every tribe that lacked firearms, no matter how far away or unthreatening. Rendering the resulting gains more permanent required that decisive outcomes—annihilation, enslavement, or permanent expulsion of the opposing tribe—be pursued wherever possible. The wars ended once all the remaining tribes were well equipped with muskets, because this meant that no more large, rapid shifts in the balance of power were to be expected. Thus, Ballara's argument, that the causes of the Musket Wars were no different than the causes of Maori warfare prior to European contact, is not supported by the evidence. The introduction and spread of muskets led to large, fleeting shifts in the balance of power that otherwise would have been extremely unlikely, given the rudimentary state of indigenous technology, and so caused wars that would not otherwise have happened.

The End of the Atomic Monopoly

The Soviet Union's successful development of nuclear weapons, less than five years after the first appearance of that technology in war, fundamentally altered the course of the nascent cold war, ending the US nuclear monopoly and resulting in forty years of nuclear-enforced stalemate between the superpowers. At least in hindsight, the case for the United States to launch a preventive war—before the Soviet Union could get nuclear weapons—seems substantial (Quester 2000). So why did it not? The answers debated in the literature focus mainly on the perceived costs of preventive war, but our model suggests an additional factor, rooted in uncertainty about how fast nuclear technology would diffuse, that might explain US hesitation.

Historians have documented serious consideration at high levels in the Truman administration of a preventive attack on the Soviet Union.[9] Two explanations are commonly given for why the United States elected not to do so. First, Buhite and Hamel (1990), Debs and Monteiro (2010), and Trachtenberg (1988, 2007) argue that, at the time, the United States perceived the material costs of a preventive war to be so high that it would rather tolerate the certain loss of its monopoly than fight to preserve it. The crux of these arguments is that, in the critical period from 1945 to 1949, the United States was ill prepared to fight a war against the Union of Soviet Socialist Republics (USSR): it had demobilized most of its military, especially in Europe where the USSR could retaliate devastatingly against a preventive attack; its nuclear stockpile and delivery capabilities were too small to cover the large Soviet target set; and it lacked the intelligence necessary to enable a more surgical preventive strike on the Soviet nuclear program. Thus, preventive war was just too expensive to be worthwhile.

These arguments raise the question of why the United States was so unprepared (Quester 2000, chap. 2). When, in the aftermath of the Soviet nuclear test and soon after the start of the Korean War, the United States commenced a military buildup, it was able greatly to strengthen its conventional forces in Europe and elsewhere and expand its nuclear arsenal in just two years (Trachtenberg 1988; Silverstone 2007). This buildup could have been implemented over 1946–48, rather than 1950–52, if the United States had chosen to do so, and would have substantially reduced the material costs of preventive war discussed above. Quester (2000) and Silverstone (2007) thus argue that the material costs of preventive war are not sufficient to explain US hesitation. Instead, they contend that the *moral* costs of preventive war, in which an unprovoked attack on the Soviet Union would result in millions of civilian deaths, explains US hesitation, because US decision makers believed that the US public would view such a war as morally unacceptable, regardless of the strategic arguments for it.

Our model, like other models of shifting power, suggests that the costs of a preventive war, whether material or moral, may not suffice to explain the occurrence or absence of such a war, and by extension, of preparations for the war. Rather, whether war occurs depends not only on its costs but also on expectations about the size, direction, and likelihood of future changes in the balance of power. Our model is unique in allowing us to analyze the incentives for war when there is uncertainty about when the balance will change.

There was widespread agreement on the part of US policy makers that the arrival of a Soviet nuclear capability would eliminate a substantial US advantage, exacerbate US conventional weakness in Europe, and expose the United States itself to the risk of a devastating surprise attack (Trachtenberg 1988; Richelson 2007; Silverstone 2007). To put it in terms of our model, the perceived difference between p_{10}, the status quo balance of power, and p_{11}, the new balance that would obtain if a shift to nuclear bipolarity occurred, was large. Proposition 2 demonstrates that this heightens the incentives for preventive war.

The United States was greatly uncertain about when the Soviet Union would acquire the bomb and was thoroughly surprised by the first Soviet test (Richelson 2007; Gordin 2009).[10] The consensus estimates of the US intelligence community, through the middle and late 1940s, asserted that the USSR would most likely test a weapon by 1953, with a low probability of a test occurring in 1952 or 1951; other influential analysts argued for much later dates. According to Gordin and Richelson, the error derived from exceedingly spotty information on the progress of the Soviet program, from ignorance of the degree of penetration of the US nuclear establishment by Soviet spies, and from a mistaken belief on the part of some analysts that the United States had managed to corner the market for fissile material.

This uncertainty can be modeled according to the cumulative diffusion model from the previous section. The United States assessed λ_t, the probability that the USSR would acquire the bomb in year t, as quite low during 1945–1950 and increasing slowly over 1951–1952, with higher probability from 1953 on. Proposition 3 suggests that the United States could simply have been putting off the costs of preventive war until λ_t was sufficiently high—that is, until the threat of Soviet acquisition was closer to imminent.[11]

Our model thus provides an explanation for why high-level US officials would advocate a massive new military build-up only *after* the first Soviet nuclear test (Trachtenberg 1988). Before the test occurred, the United States anticipated that the Soviet program was unlikely to achieve success for several more years. Because, as noted earlier, the United States could complete the necessary buildup in just two years, policy makers had every reason to procrastinate during the 1940s—the United States could count on the Soviets to avoid war until they had ended the monopoly, and could wait to rearm until this eventuality was closer. Once a "lucky" break for the USSR, based on the unexpected successes of its spies and scientists, had occurred, many US officials concluded that the United States had to be prepared for war. The high costs of such a war were nonetheless insufficient to merit continued delay in rearmament. Thus, the US hesitation to rearm was motivated not by an unwillingness to (eventually) bear the costs of a buildup but by the belief that the United States had time to spare.

Our analysis also explains why the USSR concealed its progress. If it had not, the United States could have more easily watched its program and might have attacked just before the USSR's program was successful, when the urgency of prevention would be highest. Instead, the USSR's secrecy forced the United States to rely on highly uncertain estimates of Soviet science and espionage, leading to imperfect anticipation of the timing of the shift in power. The analysis in 'Cumulative Diffusion' subsection shows that this reduces the incentives for war. The result was that the Soviets "got lucky" and were successful before the United States expected, surprising it and thereby escaping preventive attack.

To summarize, the model exposes an additional factor that may explain US hesitation in preventing the spread of nuclear weapons to the USSR and why the Soviets took care to conceal their progress.[12] The model itself cannot adjudicate between the claim that the high costs of prevention, moral and/or material, were sufficient to cause US reluctance and the claim that the low US expectations of the speed of nuclear diffusion were also necessary. But it does show that the evidence typically given for a cost-based explanation for US hesitation is inconclusive. US lack of preparation for, and willingness to air moral qualms about, a preventive war is consistent both with an assessment that the costs would be too high and with an assessment that it could afford to bide its time.

The 1967 Six-Day War

In 1967, Israel fought a short, intense war with Egypt, Syria, Jordan, and Iraq. Recent research into the history of the Six-Day War suggests that Israel's nuclear weapons program played a more important role in the origin of the war than was previously appreciated (Ginor and Remez 2007; Cohen 2007; Aronson 2009), especially in explaining the interaction between Israel and Egypt. Though these views are controversial, they enable us to demonstrate the possibility, raised by our analysis, that the situation in which neither state has acquired the new technology can be more dangerous than the one in which one state has.[13]

Aronson (2009) argues that Israel feared Egypt would attempt to use force to delay or end its nuclear program, and interpreted Egyptian troop movements into the Sinai on the eve of the war as a prelude to preventive attack. According to Ginor and Remez (2007), the Soviet Union supported Egypt and was determined to stop the Israeli nuclear program, even flying strike reconnaissance missions over Israel's nuclear facilities and feeding false information to Egypt about Israel's preparations for war in order to incite Egypt to escalate the crisis. For its part, the United States had been pressuring Israel to forego nuclear weapons, and Cohen (2007) claims that Israel feared the United States would abandon it if it made its pursuit of weapons unambiguous or otherwise provoked the Arab states. Thus, isolated and threatened, Israel launched an attack on Egypt's air forces in order to preempt what it thought was an incipient preventive war over its nuclear program.

Our model makes plain the strategic basis for Israel's fears. Both Israel and Egypt had active nuclear weapons programs in 1967. The unilateral acquisition of nuclear weapons by either would have substantially shifted the balance of power: for Israel, the weapons would compensate for its numerical inferiority and lack of strategic depth; for Egypt, they would lessen the relevance of its technological inferiority. However, Israel's program had been much more successful than Egypt's, and so Israel was much closer to getting the bomb than was Egypt. In the model's terms, both $p_{10} - p_{00}$ and $p_{01} - p_{00}$ would be large, but $\lambda_{ISR} \gg \lambda_{EGY}$. Thus, even if Egypt might derive somewhat greater advantage from unilateral possession of nuclear weapons than might Israel, the large difference in the likelihoods of acquisition meant that the expected shift in power Δp strongly favored Israel. Proposition 4 implies that Egypt would face substantial incentives to launch a preventive attack.

Despite the incentives for war when neither state had nuclear weapons, once Israel had acquired them, it did not give serious consideration to attacking Egypt to prevent its duplication of the technology (Fuhrmann and Kreps 2010). Egypt signed the Treaty on the Nonproliferation of Nuclear Weapons just a year after the war and ceased to pursue nuclear weapons, rendering such an attack pointless. But even if it had continued its program, its difficulty in securing outside assistance and the parlous state of its finances after the war meant that the per-period chance of acquisition would have been very low for a long time (Einhorn 2004). Thus, according to proposition 2, Israel would still have had little reason to attack. In short, the incentives for war were greater when neither state had nuclear weapons than when one of them had acquired them.

The model implies, and this example illustrates, a more general danger inherent in the diffusion of arms. The spread of an existing military technology to new localities can cause war even before either side has it. It is especially likely to do so when local actors have observed that the technology was highly consequential in other settings, and when differences in outside assistance or technological development make one side much likelier to get the technology first. Both properties held between Egypt and Israel in 1967.

Implications

We have analyzed and illustrated a simple model that examines how the spread of new military technologies can lead to war. Our model explains why some technologies' spread is more war-prone than others, and why in some cases even the possibility of a new technology that no local state yet possesses might cause war. It sheds new light on debates surrounding the impact of European arms transfers to precolonial populations, on the reasons for US hesitation to forcefully preserve its atomic monopoly, and on the causes of the Six-Day War. We conclude with a discussion of the implications of the theory for the global impact of military innovation.

A broader implication of our work is that competition among the major powers may impose large, indirect negative externalities on the rest of the world, even aside from the direct effects of colonization, imperialism, and intervention. The generation and use of new, consequential military technologies by technologically advanced powers creates new opportunities for arms diffusion, which may in turn disrupt the balance of power among other states and so cause preventive wars. The spread of firearms, developed to fearsome effect by the competing European powers, contributed to the causes of wars in the precolonial world, including the Musket Wars, and it may still do so today as ever more sophisticated firearms are introduced to less-developed places. And the spread of nuclear weapons, invented to end yet another competition among major powers, contributed to wars in the Middle East, including the Six-Day War, and may lead to other wars if proliferation continues.

These externalities can be exacerbated by the willing transfer, as opposed to indigenous development, of military technologies from advanced countries to others. The Maori obtained firearms from European traders, and Israel's nuclear program was accelerated—increasing the incentives for Egypt to preventively attack it—by assistance from France. The existing literature on the statistical relationship between (conventional) arms transfers and conflict provides mixed empirical results (Craft and Smaldone 2002; Suzuki 2007), and scholars note a disconnect between the theory driving these studies and the literature on war and bargaining (Krause 2004). Our model can help bridge this gap by providing testable hypotheses about the effects of such transfers. First, when the potential belligerents see a higher probability of outside intervention, the incentives for the player most likely to be harmed by the intervention to launch a preventive attack increase. For instance, the entry into negotiations over arms packages by one side may raise the probability of a shift and thereby lead to an attack by the other. Similarly, when arms transfers are occurring regularly, sharp discord between the supported player and its external patron may lead to an increased expectation that the transfers will soon be halted, encouraging the client to fight now to lock in its current advantage.

Second, the higher the impact of the anticipated transfers on the prior balance, the more risk there is of war. For instance, a situation in which one belligerent of modest technological sophistication anticipated transfers from a superpower would be more prone to preventive war than one in which the external actor was another state of modest development. Similarly, if one belligerent has become highly dependent on external support for its military, the incentives for preventive attack if this support might someday evaporate are higher than if the external support was merely ancillary.

Our analysis yields three recommendations for a powerful state that wishes to avoid causing wars with its arms transfers. First, any desired change in the quantity and quality of arms transfers should be implemented slowly over time, so that discounting lessens the commitment problem. Second, signals of future changes in arms transfers should be concealed to the extent possible, so that the party that will be harmed by a shift will not anticipate it and so will have less incentive to launch a preventive war. Third, if the other side also has a patron, alterations in the quality and quantity of its arms transfers should be matched closely, if possible, to minimize shocks to the balance of power.

Acknowledgments

The authors are grateful to Amy Catalinac, Alexandre Debs, Jeffry Frieden, Nuno Monteiro, Robert Powell, Jonathan Renshon, Kenneth Shepsle, and Jane Vaynman for their comments on earlier drafts of this article. The authors also thank the audiences at the 2009 meeting of the Peace Science Society. Harvard University, the University of California at San Diego, the University of Rochester, and Yale University, where we presented earlier versions of this article.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Notes

1. Quote taken from Trachtenberg (2007).
2. We could also give the players the option of launching a limited (i.e., non-game-ending) attack aimed at temporarily lessening the probability of diffusion to the attacked side. Whether such an attack was preferred to a decisive preventive war in equilibrium would depend only on its relative cost-effectiveness, and it would occur for the same reasons as those we identify here, so we ignore this possibility for the sake of simplicity.
3. Baliga and Sjöström (2008) and Debs and Monteiro (2010) endogenize a player's choice of whether to attempt to acquire nuclear weapons, and show that this can lead to mixed-strategy equilibria in which there is uncertainty over whether a nuclear program exists. However, this did not occur in our cases: the United States knew with certainty about the Soviet nuclear program, and Egypt knew about Israel's nuclear program, long before either program was successful. The key uncertainty was about the progress a state was making, not about whether it was trying. Our model incorporates this uncertainty, since success in acquiring the new technology is stochastic and may take many rounds of attempts.
4. We only give a proof for the case where A has the technology but not B. The analysis for this subgame is mirrored by that for the subgame where B but not A has the technology, and can be obtained by switching labels.
5. We assume that B makes the first offer in this subgame, so that A might attack B before he has a chance to receive the technology. Otherwise, B could receive the technology before A could mount an attack, even if A chose to do so immediately. Our assumption rules out this artifact of the alternating-offers protocol.
6. These assumptions simplify the exposition but do not affect the qualitative character of the results, which can easily be extended to the case of nondecreasing and/or asymptotic λ_t.
7. The results are symmetric if we make the opposite assumption.
8. We draw here on the accounts of the Musket Wars given in Belich (1996), King (1997), and Crosby (1999).

9. See Buhite and Hamel (1990), Quester (2000, chap. 4), Silverstone (2007, chap. 2), and Trachtenberg (1988, 2007).
10. We treat the US detection of the Soviet test in September 1949 as the moment at which a large shift in p occurred. The moment just after the Soviet Union had acquired the technology, but before it had weaponized and deployed it, might seem an opportune one for preventive attack, but the Soviets recognized this vulnerability and delayed the test until there was sufficient fissile material for additional bombs (Quester 2000, 133). This is consistent with our model, which assumes that the change in the balance of power occurs too quickly for a state to stop it once it has begun.
11. While Gordin (2009) endorses the lack of preparedness as the most important factor in explaining the absence of a US preventive attack (pp. 60–61), he also mentions that the "lengthy estimate" of when the Soviet Union would get the bomb "discourage[d] a preemptive strike" (p. 87).
12. The pattern of concessions between the US and USSR during this period, as described by Trachtenberg (1988, 2007), is also consistent with the model. Before it acquired nuclear weapons, the USSR was careful to avoid provoking the United States to the point of starting a war, even while it consolidated its control over Eastern Europe and even though the United States had rapidly demobilized and removed its forces from Western Europe. Once the USSR had its own nuclear weapons, however, it became much more bellicose. The clearest examples were its approval and material support for North Korea's invasion of South Korea in 1950, and its approval of China's entry into that war later in the year.
13. For alternative views on the Six-Day War, see Popp (2006), Ro'i and Morozov (2008), and Laron (2010).

References

Aronson, Shlomo. 2009. "David Ben-Gurion, Levi Eshkol and the Struggle over Dimona: A Prologue to the Six-Day War and its (Un) Anticipated Results." *Israel Affairs* 15 (2): 114–34.

Baliga, Sandeep, and Tomas Sjöström. 2008. "Strategic Ambiguity and Arms Proliferation." *Journal of Political Economy* 116 (6): 1023–57.

Ballara, Angela. 2003. *Taua: "Musket Wars", "Land Wars", or Tikanga?: Warfare in Maori Society in the Early Nineteenth Century*. Auckland, New Zealand: Penguin Books.

Belich, James. 1996. *Making Peoples: A History of the New Zealanders from Polynesian Settlements to the End of the Nineteenth Century*. Auckland, New Zealand: Penguin Press.

Buhite, Russell D., and Wm. Christopher Hamel. 1990. "War for Peace: The Question of an American Preventive War against the Soviet Union, 1945–1955." *Diplomatic History* 14 (3): 367–84.

Cohen, Avner. 2007. "Crossing the Threshold: The Untold Nuclear Dimension of the 1967 Arab-Israeli War and Its Contemporary Lessons." *Arms Control Today* 37 (5): 12–16.

Copeland, Dale C. 2000. *The Origins of Major War*. Ithaca, NY: Cornell University Press.

Craft, Cassady, and Joseph P. Smaldone. 2002. "The Arms Trade and the Incidence of Political Violence in Sub-Saharan Africa, 1967–97." *Journal of Peace Research* 39 (6): 693–710.

Crosby, Ron D. 1999. *The Musket Wars: A History of Inter-Iwi Conflict 1806–1845*. Auckland, New Zealand: Reed Books.

Debs, Alexandre, and Nuno Monteiro. 2010. "Nothing to Fear but Fear Itself? Nuclear Proliferation and Preventive War." Unpublished typescript.

Einhorn, Robert J. 2004. "Egypt: Frustrated but Still on a Non-Nuclear Course." In *The Nuclear Tipping Point*, edited by Kurt M. Campbell, Robert J. Einhorn and Mitchell B. Reiss. 43–82. Washington, DC: Brookings Institution Press.

Fearon, James D. 1995. "Rationalist Explanations for War." *International Organization* 49 (3): 379–414.

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): 831–59.

Ginor, Isabella, and Gideon Remez. 2007. *Foxbats over Dimona: The Soviets' Nuclear Gamble in the Six-Day War*. New Haven, CT: Yale University Press.

Gordin, Michael D. 2009. *Red Cloud at Dawn: Truman, Stalin, and the End of the Atomic Monopoly*. New York: Farrar, Straus, and Giroux.

Horowitz, Michael C. 2010. *The Diffusion of Military Power: Causes and Consequences for International Politics*. Princeton, NJ: Princeton University Press.

King, Michael. 1997. *Nga Iwi O Te Motu: 1000 Years of Maori History*. Auckland, New Zealand: Reed Books.

Krause, Volker. 2004. "Hazardous Weapons? Effects of Arms Transfers and Defense Pacts on Militarized Disputes, 1950–1995." *International Interactions* 30 (4): 349–71.

Laron, Guy. 2010. "Playing with Fire: The Soviet-Syrian-Israeli Triangle, 1965–1967." *Cold War History* 10 (2): 163–84.

Levy, Jack S. 1987. "Declining Power and the Preventive Motivation for War." *World Politics* 40 (1): 82–107.

Popp, Roland. 2006. "Stumbling Decidedly into the Six-Day War." *Middle East Journal* 60 (2): 281–309.

Powell, Robert. 1999. *In the Shadow of Power: States and Strategies in International Politics*. Princeton, NJ: Princeton University Press.

Powell, Robert. 2006. "War as a Commitment Problem." *International Organization* 60 (1): 169–203.

Quester, George H. 2000. *Nuclear Monopoly*. New Brunswick, NJ: Transaction Publishers.

Richelson, Jeffrey T. 2007. *Spying on the Bomb: American Nuclear Intelligence from Nazi Germany to Iran and North Korea*. New York: W. W. Norton.

Ro'i, Yaacov, and Boris Morozov, eds. 2008. *The Soviet Union and the June 1967 Six Day War*. Stanford, CA: Stanford University Press.

Rosen, Stephen Peter. 1991. *Winning the Next War: Innovation and the Modern Military*. Ithaca, NY: Cornell University Press.

Silverstone, Scott A. 2007. *Preventive War and American Democracy*. New York: Routledge.

Suzuki, Susumu. 2007. "Major Arms Imports and the Onset of Civil and Ethnic Wars in the Postcolonial World, 1956–1998: A Preliminary Reassessment." *The Social Science Journal* 44 (1): 99–111.

Taylor, A. J. P. 1954. *The Struggle for Mastery in Europe, 1848 to 1918*. Oxford, UK: Clarendon Press.

Trachtenberg, Marc. 1988. "A 'Wasting Asset': American Strategy and the Shifting Nuclear Balance, 1949–1954." *International Security* 13 (3): 5–49.

Trachtenberg, Marc. 2007. "Preventive War and US Foreign Policy." *Security Studies* 16(1): 1–31.