Arming Costs Humanity More Than War

Andrew J. Coe
Vanderbilt University

Taylor R. Dalton*
University of Southern California

Abstract

Policymakers and scholars have long focused on the costs to human welfare imposed by war, while largely neglecting the costs imposed by arming. Yet arming—the building and maintenance of states' militaries—is also costly, and unlike war, most states engage in arming most of the time. We explicitly compare the costs of arming and war. The costs of arming can be measured by existing data on national military expenditures, but it is more difficult to assess the costs of war. We develop a novel approach to estimating war costs, based on using case studies of the costs of particular wars to project total costs of other wars from data on battle deaths. We find that, in the last half-century at least, the global costs to human welfare of arming grossly exceed those of war. This suggests that the costs of arming deserve greater attention than they have so far received.

*We thank Ben Fordham and Jonathan Markowitz for comments on an earlier version.

Introduction

From the late 19th century onward, a complex of international agreements, institutions, and norms was gradually assembled to prevent war, to limit the wars that did occur, and to bring an end to ongoing wars (Goldstein, 2011; Hathaway and Shapiro, 2017). In parallel, scholars of international relations focused on war as the most important phenomenon of conflictual state behavior, studying its origins, conduct, termination, and consequences.

By contrast, both policymakers and scholars have devoted much less attention to arming: the costly diversion of valuable resources like labor and capital from the economy to the military to prepare for war. International efforts to comprehensively limit arming are rare, occurring only among the major powers in the period between the two world wars and between the superpowers at the beginning of the Cold War (Vaynman, 2022). Most states at most times are mostly free to arm as much as they like, without discouragement from the international community, in fact with encouragement by the ready availability of arms transfers. Similarly, aside from a surge of interest during the Cold War, scholars of violent conflict have largely neglected the study of arming, arms races, and arms control.

We suspect that this relative neglect of arming reflects a common sense that war is the most damaging, most dangerous behavior in the waging of political disputes and hence deserves the most attention from policymakers and scholars. After all, even a relatively quick, "clean" war is quite costly, and many wars are both long and bloody, whereas arming doesn't kill anybody or blow anything up. Even if we are prepared to equate dollars and lives, arming typically consumes only a very small proportion of a state's economy.

We show that this common sense is likely to be factually incorrect. Intuitively, even though wars are quite costly events, they are exceptionally rare; whereas arming is cheaper but happens in almost all states, almost all the time. Theoretically, it is not clear if the large but sporadic destruction of war outweighs the small but constant drain of arming.

Empirically, we find that arming's cost greatly exceeds that of war. By our estimates, from 1960 to 2008, arming has cost humanity at least three times as much as war.

Evaluating which of arming and war does more damage to human welfare requires that we estimate the total costs of each. This is readily done for arming, since there is data available on most countries' military spending for most years, but it is much harder to do for war. While data are available on the deaths suffered in battle and the military spending that occurs during wars, there is no comprehensive data on non-battle deaths due to famine or disease, non-fatal injuries, economic disruption, repayment of debt incurred to fund war, or any other costly consequences of war. Comprehensive estimates of the total cost of war are available only for a small subset of wars that is unlikely to be representative of the whole.

We therefore resort to a bounding approach. We partition the costs of war into those that are easily measured—battle deaths and wartime military spending—and those that are not. We argue that the few comprehensive cost estimates for particular wars enable us to impute a maximum ratio between the unmeasured and the measured costs of any given war. We then use this maximum ratio to compute an upper bound for the total cost of each historical war. The resulting figures are almost certainly inaccurate, but we can be confident that they are biased upward: if our figures are wrong, it is because they are too high. Thus, our conclusion that arming has cost humanity more than war is robust to the error in our measurement, because the real cost of war will be even smaller than our upper bound.

Our finding has important implications for policymakers and scholars. If arming costs so much more than war, then its neglect relative to war by the relevant policy and academic communities should be reconsidered. There might be immense welfare gains to be reaped by a serious international effort to constrain arming, even relative to the gains from the current efforts to restrain war. In turn, this effort would surely be advanced by increased scholarly study of the causes of arming and the effect of particular combinations of international agreements, institutions, and norms that could be devised to limit arming.

Although our empirical analysis can only provide an upwardly-biased measure of the cost of wars, it suggests an approach to developing unbiased estimates of the cost of wars. Every theory of war we are aware of—whether about causes, conduct, or termination—stipulates that the cost of war is an important variable, and yet we have no measures of this variable. In applying and testing these theories on particular cases of war or war generally, scholars must simply ignore cost, implicitly taking this to be uncorrelated with other variables we do measure. Developing accurate measures of the cost of war is thus essential to the empirical study of war and should be a priority for scholars of war.

To our knowledge, ours is the first attempt to characterize the total costs of war for a large span of countries and time. Previous work provides comprehensive cost estimates for a specific war, such as the US wars with Afghanistan and Iraq in the 2000s (Stiglitz and Bilmes, 2008) and the Iran-Iraq War of the 1980s (Mofid, 1990). Other research measures particular costs of war across countries and time, such as battle deaths (Lacina and Gleditsch, 2005), civilian deaths due to reduced public health (Ghobarah, Huth and Russett, 2003, 2004), war-induced capital flight (Collier, 1999), post-war demographic and economic decline (reviewed in Kugler (2017)), and retarded economic growth and progress toward the Millenium Development Goals (Gates et al., 2012). We view our work as building on these efforts in an attempt to assemble an overall picture of the costs of war.

Conceptualizing the Costs of Arming and of War

There are three challenges that must be overcome in order to compare the costs of arming and of war. The first is the challenge of commensurability: we need somehow to render all the different costs of arming and war commensurable with each other, so that we can measure them all in the same units. The difficulty here is that arming involves only the loss of material things, while war entails the loss of both things and lives. How do we measure

material things in comparison with lives? The second challenge is disentanglement: we need to decide which costs count toward arming and which costs count toward war. This is no simple matter since arming and war are causally entangled with each other: arming can cause war, and war can cause arming. Should the cost of arming by a belligerent during a war be counted as a cost of arming, a cost of war, or both? The final challenge is data availability: much of the information we need on arming and wars is not available, so that we must find some principled way to impute the missing data. In this section, we explain our general approach to overcoming these three challenges. Subsequent sections turn to describing in detail our measurement of the costs of arming and of war and how we estimate these from the available data.

Commensurability

Our approach to the challenge of commensurability begins with the observation that arming and war are both costly in the economic sense of inefficiency (Fearon, 1995; Coe and Vaynman, 2020). Both involve the diversion of resources from civilian activities that would create value to military uses that can only re-allocate (through coercion) or reduce (through destruction) value. In arming, governments spend money on their militaries that could instead be used to fund social welfare programs or scientific research, or to reduce taxes or borrowing. Factories and skilled labor that firms devote to manufacturing fighter jets and tanks will not make airliners and cars. Time that people spend in the armed services cannot be used to produce goods and services in the civilian economy. In war, people are injured or killed and matériel is consumed, damaged, or destroyed. These people and matériel therefore suffer a permanent loss of well-being or of value, respectively. These losses are incurred directly, as a result of fighting, and indirectly, because of war-induced disruption to the economy, government services, and public health.

Conceptually, the cost of diverting some resource to military use is equal to the value

it would have created if it were instead left to civilian activity. Theoretically, this can be measured as the fair market price of the resource in question: that is, the price the resource would fetch in the economy from which it is drawn. Market prices therefore create a basis for comparing the costs of different resources, typically expressed in local, contemporaneous currency units which can be adjusted for inflation and exchange rates to a single, constant unit of value across time and countries. We use this approach throughout the paper to overcome the challenge of commensurability.

This approach is uncontroversial when it comes to resources like the funds used to finance a military, the goods and services used to equip it, and the labor-hours used to staff it. The value of these resources is conventionally measured in terms of interest rates, prices, and wages. These suffice to enable us to measure a total cost of arming. By contrast, the total cost of war involves not only things we usually appraise according to the market, but also things we do not, such as human injury and death. To make these costs commensurable, we need to know: how many dollars is a person's life worth? How many injuries are needed to equal the cost of a death?

One might object that the value of a life cannot ethically be equated to any finite amount of dollars. This objection is fundamentally misplaced. In a world of limited resources, both individuals and governments routinely make decisions that imply a finite dollar value of human life. Workers evaluate whether the pay offered for a job is worth the risk of injury or death it entails, and often take more dangerous jobs in exchange for modest increases in compensation. Consumers are willing to pay more for safer products (e.g., cars with better crash-test results), but only somewhat more. For their part, governments have no choice but to trade off dollars and lives when setting a variety of policies, ranging from speed limits to pollutant restrictions to sanitation standards. Whatever its philosophical merits, as an empirical matter, governments and individuals routinely act as though a human life has a finite dollar value.

Economists have developed statistical measures of the value of human life implied by workers' demand for extra compensation for higher-risk jobs and consumers' willingness to pay more for safer products: the result is called a "value of statistical life" or VSL.[1] VSL has been measured across many countries and over time. Moreover, national governments increasingly set policies according to standardized monetary values of their citizens' lives, which are drawn from economists' measures of average individual VSLs. Because there is increasing convergence around measuring the value of human life by VSL, we adopt it here. This enables us to measure a total cost of war and compare it to a total cost of arming expressed in the same units.

Measures of VSL vary substantially across countries and time. This variation corresponds closely with the contemporaneous gross domestic product per capita (GDPPC) of the country for which VSL is measured. Individuals and governments in countries with higher income per capita tend to act as though they value a human life more highly (in dollar terms) than individuals in countries with lower income per capita, on average. This relationship is approximately linear, so that VSL in a given country at a given time is approximately 150 times the GDPPC of that country at that time. We use this relationship to value each human life lost at 150 times the contemporaneous GDPPC of that person's home country, drawn from the World Bank.[2].

This variation itself might seem objectionable: is a life in one place, at one time, really worth fewer dollars than a life at some other place or time? But the same reasoning

[1]For a thorough review of the intellectual history of and scholarly literature on VSL by the concept's originator, see Viscusi (2018).

[2]Of course, an ongoing conflict that causes deaths in a particular country may also disrupt that country's economy, depressing its GDPPC. If GDPPC would be higher in the conflict's absence, then using the contemporaneous war-depressed GDPPC might cause us to undervalue battle-deaths. We analyzed this possibility for a randomly selected subset of conflicts, and tested alternatives such as using the GDPPC observed just before a war starts, the GDPPC observed just after the war ends, and the highest GDPPC observed within the five years before and after the war to value lives lost during the war. We found that these alternative choices made very little difference in the total valuation of lives lost in our test subset. As a result, we chose to stick with the simpler approach of using contemporaneous GDPPC.

that justifies the use of VSL also justifies incorporating its variation. Individuals and governments everywhere face a world of limited resources, but some places and times are less resource-constrained than others. When and where resources are more plentiful, individuals and governments behave as though lives are worth more. Whatever moral quandaries this imposes, our approach values human lives according to the valuation implicitly used by people living in the same place and time.

Disentanglement

Using VSL and market prices enables us to render all the different costs of arming and of war commensurable, but still leaves us with the challenge of how to disentangle arming from war. At first blush this might seem straightforward: arming entails the costs associated with preparing for war, while war involves the costs of its actual conduct. But then how should we count military spending during a war? Some of this is incurred in waging the war, but some would probably have occurred even in the war's absence. What if the mounting costs of an ongoing arms race lead one side to launch a war against the other—should that be counted as a cost of war or of arming? Because the extent of arming can affect whether war occurs, and the occurrence of war can affect the extent of arming, we need to somehow divide the costs of the two. We proceed to enumerate the known causal mechanisms linking arming and war, argue that all but one of these can be conservatively ignored, and explain our approach to dealing with the one exception.

First consider how unilateral changes in arms affect war. If one state unilaterally increases its arms, this could reduce the occurrence of war by strengthening deterrence, or instead increase war by emboldening the more-armed state. The bargaining theory of war holds that these two effects should balance each other out: the more-armed state is emboldened exactly to the degree that its opponents are deterred. Whether war occurs or not depends only on whether changes in arming affect bargaining failures through commitment problems or

asymmetric information (Fearon, 1995). On the one hand, if there is a delay between seeking arms and deploying more arms, then changes in arming could increase war by encouraging the soon-to-be-weaker state to launch a preventive attack (Debs and Monteiro, 2014). A state's change in arms could also encourage war by generating uncertainty on the part of an opponent as to how much the state's military power has been altered (Meirowitz and Sartori, 2008). On the other hand, if an increase in arms is motivated by the need to match an anticipated increase in an adversary's strength, then this arming averts the shift in the balance of power that would otherwise occur, thereby eliminating a commitment problem and discouraging preventive attack. And an increase in arms could just as easily dispel as create uncertainty about the balance of power, if it increases confidence that a state's military is ready and able to fight effectively.

We can avoid the indeterminacy of the effect of unilateral arming on war by instead focusing on *overall* arming. Because we seek to evaluate the global costs of arming and of war, and the implied relative benefits of global efforts to reduce arming or war, it is reasonable to consider only the effect of overall arming. On this, existing theory is clear and determinate. In principle, overall arms can be increased or decreased with no effect on the balance of power between two sides, as one side's change in arms can be made to match the other's (Coe and Vaynman, 2020). But increased overall arming raises the cost of the two sides remaining at peace with each other—heightening the level of arming is expensive—and this increases the incentives for war as a way to avoid the peacetime costs of arming (Coe, 2011; Fearon, 2018; Powell, 1993). It also decreases the marginal cost of going to war, since more of the arming needed to wage war has already been paid for, again encouraging war (Slantchev, 2005). Thus, higher overall arming should increase the occurrence of war, meaning that some portion of the overall cost of war should be attributed to the overall cost of arming.

The same mechanism also works in reverse. If a war leads to the decisive defeat of one

side, eliminating or at least durably lessening the future threat it will pose, then both sides' subsequent arming will be reduced and the overall level of arming will thereby decline. Put another way, to the extent that war occurs to avoid further peacetime costs of arming, then the war's assessed cost should be reduced to account for the future savings in arming the war will induce.

Finally, even if war would decrease future arming, it may increase contemporaneous arming as the belligerents devote additional resources to their militaries in order to wage the war. This is commonly observed historically and clearly constitutes a substantial portion of the arming that we observe in states at war. We will call this "war-induced arming" in order to distinguish it from arming that would have occurred even if the war had not.

Collecting the mechanisms: overall arming increases war, while war decreases future arming but increases contemporaneous arming. We will henceforth ignore the first two, operating as though arming has no impact on the subsequent occurrence of war and war does not affect future arming. Because this choice will bias our estimate of the total cost of arming down and our estimate of the total cost of war up, it is conservative. Our headline conclusion is that the global cost of arming exceeds that of war even ignoring these two mechanisms: incorporating either would only strengthen that conclusion.

To deal with the last mechanism, we will compute two versions of our estimates of the costs of arming and of war. In one version, all arming is included only in the total cost of arming, even war-induced arming by belligerents waging war. In the other version, we partition wartime arming into a portion that was induced by war and a remainder that would have occurred even in the war's absence. The former is counted only toward the cost of war, and the latter is counted only toward the cost of arming. This version accounts for the possibility that much of the observed arming would not have occurred but for ongoing wars, and moves the cost of that arming from the arming to the war side of the ledger.

Which version of our figures is relevant depends on the policy instrument that would be

employed to lower the overall cost of arming. Consider the possibility of a new multilateral arms control agreement that is adopted in order to lower arms levels. If it applies (and is complied with) in all situations, then it can lower all arming, even war-induced arming, and so the version of our estimates that counts all arming toward the cost of arming is appropriate. If instead the agreement will only be obeyed when signatory states are not at war, then it can only lower arming that is not war-induced, so that the version of our figures that removes war-induced arming from the cost of arming applies. There are historical examples of both. The interwar treaties on naval forces were generally complied with in peacetime and contained no wartime exceptions, but were abandoned by all parties in World War II, while the nuclear nonproliferation treaty generally has been obeyed even when participants are at war. We therefore report both versions of our estimates of the costs of arming and of war.

Data Availability

Even given an approach for rendering all costs commensurate and for allocating them between arming and war, we are missing much of the information we need to estimate these costs. This last challenge applies to both arming and war. For arming, the available cross-national, over-time data is limited to total military spending and personnel, and even that data is incomplete, with missing values for some country-years. For war, there are data on battle-deaths across conflicts and years, though with a considerable portion of conflict-years missing a value. We have very little specific data on all the other costs of war, for most wars. In fact, we are aware of only four comprehensive estimates of the total costs of war for particular states: the US in the US-Afghanistan War and in the Iraq War (Stiglitz and Bilmes, 2008), and Iran and Iraq in the Iran-Iraq War (Mofid, 1990).

This challenge seems dire, but it is not insurmountable. For each component cost of arming or war that we must measure, the method we use depends on how much information

we have about the missing values. If only a modest fraction of the values of some component cost is missing, we utilize the Amelia II R package to fill in multiple "guesses" for the set of missing values (Honaker and King, 2010). All of our qualitative conclusions are robust across these guesses, as well as to the use of simpler single-imputation techniques including linear interpolation and using the mean of observed values to fill in missing ones. This works well for military spending (5.5% missing from 1960 to 2012), gross domestic product per capita (13% missing), and the best estimate of battle deaths (38% missing), where only a modest share of state-years (or conflict-years) have missing values.[3]

For other variables, most of the values are missing. This makes statistical imputation of particular values less reliable, because there is relatively little information available about the distribution of these values. We therefore try to determine the feasible range of each missing value, rather than guessing a particular point value. The construction of these ranges is necessarily tailored to each particular variable, and will be described subsequently. We will show that our qualitative conclusions are robust to any choice of values within the feasible ranges, so that the remaining uncertainty about the true value can be ignored.

Finally, when it proves impossible to specify a feasible range, we resort to finding only a conservative bound on each missing value. The main conclusion of this paper is that, over the half-century from 1960 to 2008, the global cost of arming exceeded that of war. Hence, "conservative" means finding a lower bound on any variable that adds to the cost of arming, and an upper bound on any variable that adds to the cost of war. This has the downside that our estimates of the total costs of arming and of war are not unbiased: the figure we compute for the cost of arming is surely too low, and that for the cost of war is very likely too high. However, if even given these biases, the estimated cost of arming exceeds that of

[3]The missingness in battle deaths is not as bad as it appears. Although we are missing more than a third of values for the point estimate of battle deaths, the dataset we use supplies lower and upper bounds on battle deaths for each of these missing values, which reduce the uncertainty considerably and are incorporated into the imputation procedure.

war, then we can be confident that the true (higher) cost of arming is larger than the true (lower) cost of war. Thus, the upside of our approach is that it enables us to show that the cost of arming is greater than that of war despite the severe limitations in the available data. Because our goal is only to determine which of arming or war costs more overall, even relatively crude measures of the costs of both will do.

Measuring the Costs of Arming

We proceed to describe our approach to measuring the cost of arming using military spending, and argue that this is a conservative (i.e., too-low) estimate of the cost of arming. We then explain how we estimate the degree of arming that is war-induced, so that this can be removed from our figure for the total cost of arming and added to that for war.

Voluntary versus Coercive Arming

Arming requires governments to divert resources from the civilian economy to staff, equip, and fund their militaries. Governments have two means to accomplish this diversion: inducing the owners of the desired resources to voluntarily agree to diversion, or coercing them into doing so by the threat of force. As we will explain, the mix of these two means chosen by a given government at a given time affects how we should go about measuring the costs of arming.

Voluntary diversion requires a government to pay soldiers enough that some people prefer military service to working in the civilian economy, to pay enough for military equipment that some firms can profitably produce it, and to offer interest rates high enough for some investors to prefer buying government debt, the receipts of which can be used to fund the military, to investments in the private sector. In economic equilibrium, a government should pay wages, prices, and interest rates that are just high enough to render private owners of

desired resources willing to volunteer them for military use. If it pays more, it can reduce the compensation it offers while still obtaining the desired resource. Thus, the compensation offered for military service should be no higher than the value the diverted resources would fetch in the civilian economy: their fair market price. This means that we can use a country's military spending to measure the cost of all resources voluntarily diverted to military use.

Of course, governments do not always offer full compensation for the value of resources used in the military. When the compensation offered is lower than the civilian-economy value for the resources, governments must use coercion to divert these resources. Coercive diversion involves a government conscripting citizens into military service, commandeering firms' productive capacity for military use, and taxing the civilian economy to finance the military budget. Though governments do not directly pay for it, coercive diversion is still costly, because the civilian economy still loses the value that would have been produced by the resources coerced into military service. However, we are aware of no dataset that records the extent of coercive diversion across countries and over time.

We therefore conservatively adopt a country's military spending as a lower bound on its cost of arming. This spending is a lower bound because, while it includes the full cost of all resources voluntarily diverted to the military, it excludes the cost of any resources diverted through coercion. This choice of measure is conservative because, having found that the measured cost of arming is lower than the cost of war, we can be all the more confident that the true cost of arming, which is actually higher than our lower-bound measure, exceeds that of war. We draw annual military spending figures for most countries in most years from the National Military Capabilities dataset, version 5.0 (Singer, Bremer and Stuckey, 1972; Singer, 1987), and use statistical multiple-imputation to filling in the missing country-years.

To assess the robustness of our subsequent findings, it is important to note that this lower bound is certain to be well below the true cost of arming. All arming entails at least some positive level of coercive diversion. Even if a government fully compensates

soldiers, firms, and investors, the interest it offers investors in order to obtain the funds to compensate soldiers and firms must still be paid for through taxation, a coercive form of diversion.[4] Taxation withdraws capital from the civilian economy which could otherwise fund productive investment, and so imposes costs. These costs can be quite substantial, because capital generates high returns in the economies of many countries that have high military spending, especially Western countries like the United States and United Kingdom.

We ignore any positive externalities of military spending. Spending on military research and development (R&D), for instance, sometimes results in innovations that create value in the civilian economy. This added value of military spending should be subtracted from its total cost to arrive at its net cost to the economy. However, even in the United States, which devotes far more funding to defense R&D than any other country over the period we study, R&D is less than one-eighth of the typical year's military spending.[5] Moreover, economic studies of the externalities of military R&D tend to find that these are small: valuable civilian spin-offs like the internet are more the exception than the rule (Martí Sempere, 2018; Moretti, Steinwender and Van Reenen, 2019).

War-Induced Arming

As explained in a previous section, we will produce two estimates of costs of arming and of war. In the first, we ignore all causal relationships between arming and war, treating the cost of arming as excluding all costs of war, and the cost of war as excluding all costs of arming. In the second estimate, we conservatively ignore all causal relationships between arming and war except for the one by which the occurrence of war induces belligerents to increase arming in order to wage the war. We remove the costs of this war-induced arming

[4]Governments can instead print money, but this generates inflation, which is an indirect tax on wealth.

[5]For defense R&D budgets across countries, see the indicator "Government budget allocations for R&D" in the Research and Development Statistics of Organisation for Economic Co-operation and Development (2022). We calculated the share of R&D in annual US defense budgets from Table 3.2 in Office of Management and Budget (2019).

from our total cost of arming and add it to the total cost of war. Here, we explain how we estimate how much of the arming done by belligerents during a war is in fact war-induced, as opposed to arming that would have occurred even had the war been absent.

Our approach is designed to be conservative, in the sense that is almost surely overestimates how much arming is war-induced. We classify each conflict as either "existential" or "limited" for each involved state. When a conflict is coded as existential for a belligerent, we assume that all of its wartime arming is war-induced and thus counted toward the cost of war, not that of arming. When a conflict is coded as limited for a belligerent, only a fraction of its wartime arming is assumed to be war-induced: we compute results for 5, 20, or 40 percent.

Any conflict that causes less than 1000 battle deaths is coded as limited, except for a small state fighting on its own territory against a great power. Among conflicts that exceed this threshold, internal conflicts are still coded as limited if they are not center-seeking. Interstate conflicts are limited for a great power unless the opponent is also a great power and the conflict takes place within a great power's territory. They are limited for other states unless the conflict takes place within the state's territory. All data on conflicts is drawn from the UCDP/PRIO Armed Conflict Dataset (Gleditsch et al., 2002; Allansson, Melander and Themnér, 2017).

These rules result in a classification of conflicts that corresponds closely with our intuitions about which wars are limited and which decidedly less so. To the extent that they err, they do so in the direction of overclassifying conflicts as existential and are thus conservative.

Measuring the Costs of War

The central challenge of adding up the costs of war is that we have available comprehensive estimates of the total cost of war for only four belligerent states across three wars: the US

in the Afghanistan and Iraq Wars and Iran and Iraq in the Iran-Iraq War. For all other conflicts, we know only one component of the cost of war: the estimated number of battle deaths. Here we describe our approach to overcoming this challenge. We will argue that the few comprehensive estimates we have should be outliers in a certain sense relative to other conflicts, and that these outliers can be used to construct an upper bound on the cost of those other conflicts.

War as Economic Activity

We begin by conceptualizing war as an economic activity in which the belligerents expend their citizens' lives in battle as well as other resources from their economies in order to (attempt to) produce their desired outcome for the war. We can then partition the total cost of war into two categories. The first is the cost resulting from battle-deaths, valued using the VSL methodology discussed previously: $C_{cb}(\text{battle-deaths}) = VSL_{cb} * BD_{cb}$, where VSL_{cb} is the VSL for belligerent b in conflict c and BD_{cb} is that state's number of battle-deaths in that conflict. The second category is the cost of all the other resources expended in the war. This includes non-battlefield deaths and non-lethal injuries to people; damage to and destruction of physical assets such as buildings, infrastructure, productive land; and disruption of the belligerents' peacetime economies. Given this conceptualization, we can write the total cost of each conflict c for each belligerent b as:

$$\begin{aligned} TC_{cb} &= C_{cb}(\text{battle-deaths}) + C_{cb}(\text{other resources lost}) \\ &= VSL_{cb} * BD_{cb} + C_{cb}(\text{other resources lost}) \\ &= VSL_{cb} * BD_{cb} * \alpha_{cb}, \end{aligned}$$

where $\alpha_{cb} \equiv \frac{C_{cb}(\text{battle-deaths}) + C_{cb}(\text{other resources lost})}{C_{cb}(\text{battle-deaths})}$ is a multiplier that is specific to each conflict-belligerent.

Because we know the number of battle-deaths for every conflict, our problem is now reduced to estimating the value of the multiplier α_{cb} for each conflict-belligerent that lacks a comprehensive cost estimate. Our key theoretical claim, upon which the rest of our measurement approach depends, is that this multiplier is tightly coupled to the belligerent's level of economic development.

Economic activity entails the use of inputs like labor and capital to produce valuable goods and services. Development increases the supply of capital relative to labor and so leads to an increasing substitution of relatively more abundant (and therefore cheaper) capital for relatively scarcer (and therefore costlier) labor in production. For example, farming in poorer countries typically involves gathering crops by hand and is thus labor-intensive; in richer countries, many field hands are often replaced by sophisticated harvesting machines operated by a single, more skilled worker and is instead capital-intensive. Even sectors which are relatively labor-intensive in all economies, such as education, still feature higher capital-to-labor ratios in richer countries: computers are used and exams are sometimes graded by scantron rather than by the teacher. Development thus increases the ratio of capital to labor in the economy generally and also in most specific sectors of the economy.

Militaries are no exception to this general phenomenon. In richer countries' armies, soldiers are given more and higher-quality training as well as more and better equipment. Their navies and air forces use more technologically-advanced warships and aircraft with smaller crews. Their operations rely more heavily on sophisticated communications and intelligence systems. All of these represent substitutions of capital for labor. Because militaries must ultimately draw their inputs from the broader civilian economy, they too must find ways to replace costlier labor with cheaper capital, so that the ratio of capital to labor in militaries is driven by that of the broader economy.

Violent conflict damages and destroys capital and labor within the involved militaries but also within the broader civilian societies of the belligerents, especially if the conflict

takes place in their own territory. In military forces, capital and labor tend to come tightly packaged, so that labor and capital tend to suffer together. Sinking a warship, shooting down a combat aircraft, or destroying an armored vehicle all involve the destruction of both capital (the vehicle and its equipment) and labor (its crew). Soldiers attacked on the battlefield may lose their lives along with their equipment. Similarly, in the broader civilian society, aerial bombing and artillery strikes that target factories, homes, and infrastructure will also harm the people working and residing in those facilities. Blockades, sanctions, or attacks on commerce disrupt the economy and may lead to disuse of both workers and physical plant.

Thus, the capital to labor ratio of resources expended in conflict by a belligerent should be strongly dependent on their ratios in the belligerent's military and in its broader economy, and the ratio in each state's military should also be closely correlated with that of its economy. We can demonstrate the latter claim by simply examining the capital-to-labor ratio of states' militaries. Figure 1 displays, for all states from 1960 to 2012, the relationship between development (measured by gross domestic product per capita, GDPPC) and military capital-to-labor ratio (CLR, measured by annual military expenditure divided by military personnel). The solid and dotted curves are loess curves fitted to wartime and peacetime state-years, respectively. Clearly, development and the military CLR are tightly coupled, with the relationship between them approximately linear (with both variables logged). A simple regression of military CLR on GDPPC explains 23% of the variance in the former. This gives us confidence that the same coupling should be present in war, though it is important to point out that we cannot directly test this claim.

Battle-deaths are the costliest form of damage to labor in violent conflict. Using labor in military operations is costly even if no one is harmed, because the time that each person spends in the military is lost to the civilian economy, costing the state whatever value in goods and services the person could otherwise have produced. However, deaths from military

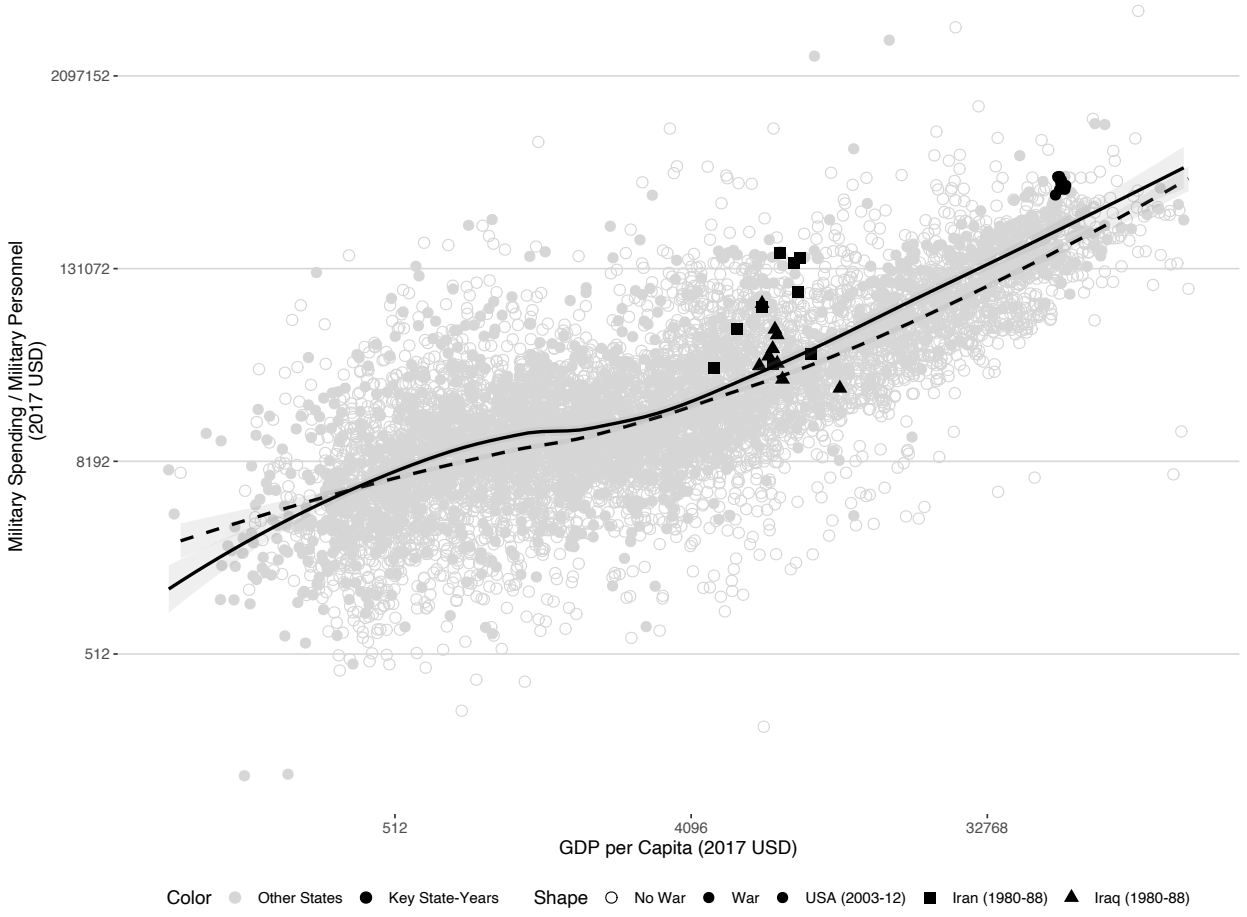

Figure 1: The Relationship between Military Capital-to-Labor Ratio and Development

operations entail the permanent loss of each person's production, not just a temporary loss while they serve in the military. More developed states thus have the strongest incentive to reduce military deaths, even relative to other expenditures of labor, by the employment of more capital: soldiers wear body armor; tanks feature active armor; aircraft are equipped with ejection seats; and medical evacuation and treatment are generously funded.

[Should also have graph of battle-deaths per military expenditures by GDPPC, showing tight correlation and explaining that this is evidence that more developed states substitute capital losses for labor losses in warfare just as they substitute capital for labor in their militaries and economies.]

Returning to our decomposition of the total cost of conflict, the multiplier α_{cb} is equal to the ratio of all the expended resources—lives lost and all others—to battle-deaths in a particular conflict by a particular belligerent. Because this ratio should be tightly coupled to a state's contemporary level of development, the multiplier can be estimated as a function of development: $\alpha_{cb} = f(GDPPC_b)$. We can use the few available comprehensive estimates of war costs for particular conflict-belligerents to calibrate that function, and then use the calibrated function to predict each belligerent's multiplier in each conflict based on its GDPPC. We turn next to this calibration.

Calibrating the War-Cost Multiplier as a Function of Development

In an ideal world, we would have comprehensive estimates of the total cost of war for a sample of belligerent-conflicts that spanned the observed range of economic development and was representative of the population of belligerent-conflicts. We could then calculate the value of the war-cost multiplier for each belligerent-conflict in the sample by dividing the total cost of war for each by its battle-deaths, and plot these multipliers against the belligerents' GDPPC. This would enable us to infer the shape of the function relating development to the war-cost multiplier with some confidence and to make unbiased predictions of the values

of the multiplier in all the cases where we don't know the total cost of war, but do know the belligerent's GDPPC.

The world is not ideal. First, the cases we have comprehensive estimates for do not span the observed range of development: the US during the Afghanistan and Iraq Wars is near the upper extreme of GDPPC, while Iran and Iraq in their war had similar GDPPC that are much lower than that of the US but still well above most state-years in the data. We thus have big gaps in coverage for relatively poorer states and for those states with development higher than Iran/Iraq but lower than the US.

Second, there is good reason to think these four cases are unrepresentative even among belligerents with similar levels of development. The US has exceptional military technology even relative to similarly rich countries, owing to decades of large investments in military research and development (Brooks and Wohlforth, 2015). As we would therefore expect, Figure 1 shows that its military has a capital-to-labor ratio that is higher than that of most other states of comparable (or lower) GDPPC. Similarly, at the time of their war, Iran and Iraq were widely viewed as possessing two of the most technologically-advanced militaries among developing states. Thanks to its prior client relationship with the US (which ended just before the war), Iran fielded the best military equipment the US exported and its forces had received extensive training from the US. For its part, Iraq had received analogous equipment and training from the Soviet Union. We can see from Figure 1 that both Iran and Iraq are also outliers in military capital-to-labor ratio, relative to states of comparable or lesser GDPPC.[6]. We can thus infer that, relative to states of comparable or lesser development, the US, Iran, and Iraq should each have unusually high multipliers.

Because the few cases for which we have comprehensive estimates are concentrated at just two levels of development, we cannot pin down the shape of the function relating GDPPC

[6]Iraq's CLR ranges widely due to the rapid and very large expansion of its army during the war, but the average over the war is still well above the mean for similarly-developed countries.

to the war-cost multiplier with confidence. We have made a theoretical argument that the function is increasing in GDPPC, and we know it must intersect with the multiplier values of our comprehensively-estimated cases, but we do not know how fast it increases or how its slope varies over the observed range of GDPPC. Even if we assume a shape to the function, because those same cases are unrepresentative, our multiplier predictions will be biased.

However, all is not lost. Because we are confident that the comprehensively-estimated cases should have unusually high multipliers for their levels of development, we can use them to place upper bounds on the distribution of multipliers as a function of development. These bounds then enable us to estimate an aggregate upper bound on the total cost of war across belligerents and conflict. Because the figure we derive will be higher than the actual total cost of war, and yet still lower than the total cost of arming we find, we can be confident in our conclusion that the cost of arming is greater than that of war.

We impute the unknown multipliers according to two methods, each based on a different assumption about the shape of the function relating the distribution of multipliers to development. In the first method, we assume the function relating GDDPC to multiplier has the same shape as the function that relates GDPPC to military capital-to-labor ratios, which Figure 1 shows is approximately linear. We thus fit a linear model to the known multipliers and GDPPC of our comprehensively-estimated cases and then use this model to predict multipliers for other cases based on their GDPPC. Given the assumption of linearity, this method produces conservative (because too-high) multipliers for other belligerent-conflicts on the basis of the observed slope between the (unusually high) multipliers for the US, Iran, and Iraq.

Our second method is even more conservative and eliminates the need to assume anything about the shape of the function relating development to multipliers, other than that it is increasing, as we have argued. We simply assume that the multiplier is a step function of development. This means that every belligerent that is less-developed than Iran and Iraq

(which have similar GDPPC at the time of their war) is imputed to have the same multiplier they do (it will turn out the two have roughly the same multiplier from the comprehensive estimates). Every belligerent that is more-developed than Iran and Iraq but less-developed than the US is imputed to have the same multiplier as the US.[7] This results in predicted multipliers for each belligerent-conflict which are the highest possible that are still consistent with the multiplier being an increasing function of development, regardless of the shape of that function. It is therefore extraordinarily conservative, as the multipliers imputed via this method will always be larger than those of the first method, and often much larger.

Figure 2 illustrates the two different methods of imputation. The light blue dots represent the (unknown) values of the multipliers for belligerent-conflicts for which we do not have comprehensive estimates of the total cost of war. The dark blue dots represent the known multipliers for the US in the Afghanistan Wars (the upper-right dot) and for Iraq and Iran in their war (the lower-left dot). In accordance with our theoretical argument that the US and Iran/Iraq should have higher multipliers than most other belligerents at or below their respective levels of development, the two dark blue dots are higher than most points to their left. The red line represents the linear model based on these known multipliers, from which we impute all the unknown multipliers. We have placed the unknown multipliers (light blue dots) as though they are linearly related to development, just like we showed the military capital-to-labor ratios to be. However, we do not know this to be true. The green lines represent our more conservative method, where we impute the unknown multipliers from a step function anchored by the known multipliers. In this second method, each imputed multiplier is as large as it can possibly be without violating the much weaker assumption that multipliers generally increase in development.

[7] We assume the belligerents with higher development than the US also have the same multipier as the US, because as we explain below they only appear in conflicts in coalition with the US, and suffer a negligibly small share of the coalition's battle-deaths.

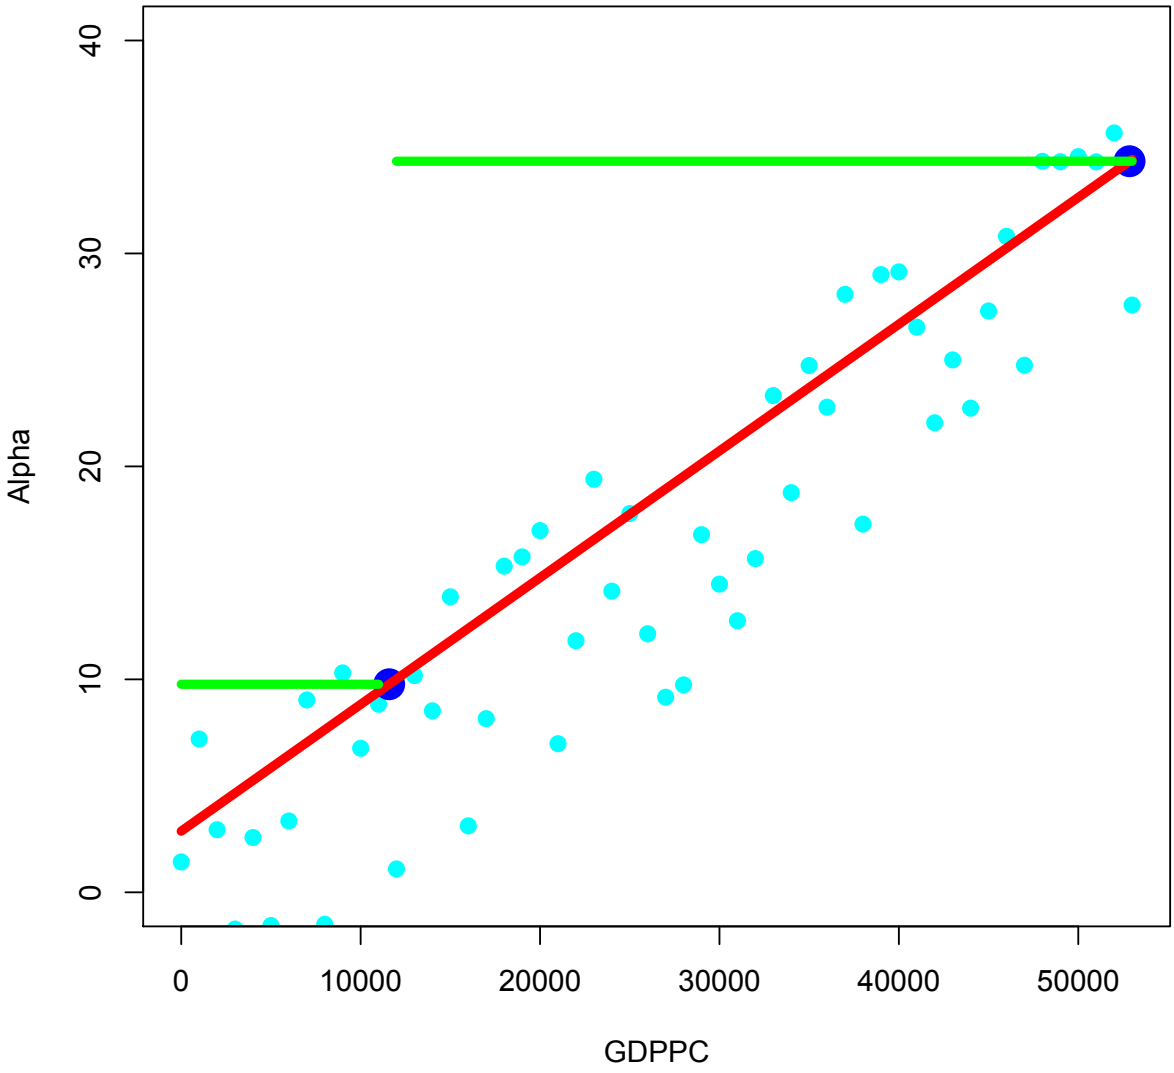

Figure 2: Two Methods for Imputing Unknown War-Cost Multipliers from Development

Partitioning Battle-Deaths among Belligerents

Our method of converting human lives to dollars requires that we know not only how many died in a war, but also how these deaths are partitioned across the belligerent states, so that we know which country's GDPPC to use to calculate the value of statistical life for each death. The battle-death data we use (Lacina and Gleditsch, 2005, Version 3.0) does not provide this information, so we instead compute the range of feasible values. At the lower end of the range, we assume that all battle-deaths were suffered by the belligerent with the lowest GDPPC; at the upper end, we assume all the dead came from the highest-GDPPC belligerent. In cases where the difference between the richest and poorest belligerents is large, the resulting uncertainty is also large. To reduce this uncertainty in our estimates of the value of battle-deaths, we used a variety of sources to fill in country-specific battle-death figures for every belligerent with a contemporaneous GDPPC above $10,000 in 2017 US dollars.

In large-coalition conflicts such as the US-led invasions of Afghanistan and Iraq, we often only have figures for all coalition deaths, not figures specific to particular coalition members. We resort to attributing all coalition deaths to the US. Because the US military provides most of the coalition forces and suffers most of the casualties, this is a reasonable approximation. However, it is also conservative because the US is typically among the richest members of the coalition, so that our approximation will tend to bias our estimate of the cost of battle-deaths upward, militating against our conclusion that war costs less than arming.

Results

Table 1 reports the values we calculated for the multipliers for our four comprehensively-estimated cases of the cost of war for particular belligerents: the United States in the Afghanistan and Iraq Wars, and Iran and Iraq in the Iran-Iraq War. (We are still work-

Table 1: Multipliers for Comprehensively-Estimated Cases

War-Induced Arming Cost	Iran	Iraq	US in Afghanistan	US in Iraq
excluded		9.77		34.3
included		10.9		62.5

ing on the Iran and US-Afghanistan calculations. Details of these calculations appear in the online appendix.) The multipliers are calculated in two different ways, depending on how war-induced arming is treated. In the top row, all military spending during each conflict is excluded from the total cost of war. In the bottom row, military spending that was initiated in order to fund war efforts is included in the total cost of the conflict; hence, these multipliers are higher.

We then use these multipliers as references from which to impute multipliers for all other belligerent-conflicts in the data, according to the level of development of each other belligerent, by linear model and by step-function as described in the previous section. Taking the product of the imputed multiplier and the battle-deaths, valued according to the belligerent's VSL, we arrive at our estimate of the total cost of each conflict for each belligerent. Summing these costs across belligerents and conflicts in each year from 1960 to 2008, we get the annual total cost of conflict shown in Figure 3.

Starting from the bottom of the figure, the orange and red lines represent the valuations of lives lost in battle during conflicts in each year. The gap between the two lines arises from our uncertainty in some cases about how to allocate battle-deaths among belligerents, but this uncertainty is modest in most years. We can see pronounced bumps in the late 1960s, heavily driven by US and Vietnamese battle-deaths in the Vietnam War, and a larger plateau for most of the 1980s, driven by Soviet and Afghan deaths in their war and Iran's and Iraq's in theirs. Moving up, the purple and magenta lines represent the total cost of conflicts for all belligerents in each year, where the multiplier for each conflict-belligerent is linearly interpolated from the known US and Iraq values. Here too, the late 60s and the

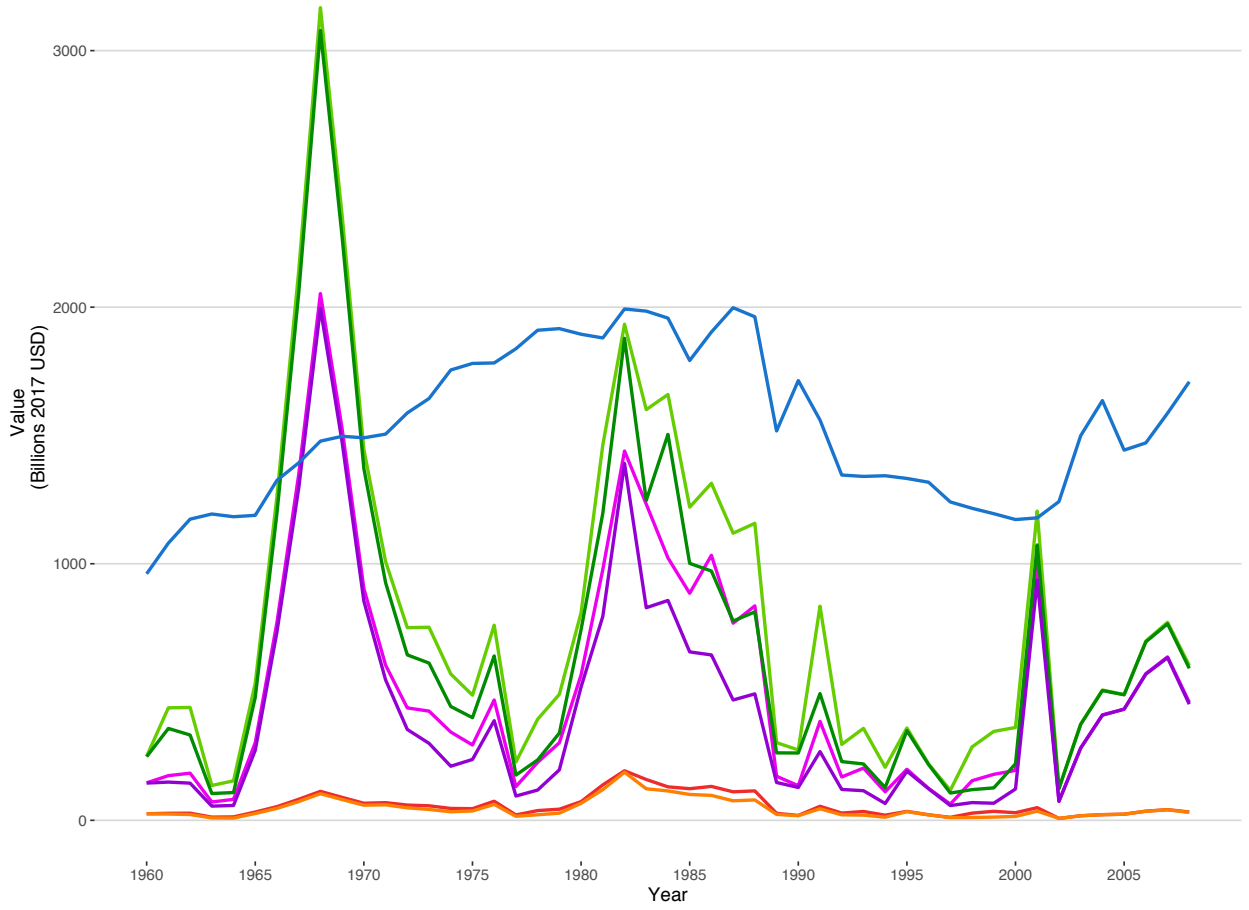

Figure 3: The Annual Global Cost of Battle-Deaths (orange to red lines), Conflict by Interpolation (purple to magenta lines), Conflict by Step Function (dark to light green lines), and Arming (blue line)

Table 2: Ratios of Total Cost of Arming to Total Cost of Conflict

Method of Multiplier Imputation	Wartime Military Spending Assumed to Be War-Induced			
	0%	5%	20%	40%
Linear	2.95 - 3.49	2.25 - 2.57	2.10 - 2.40	1.91 - 2.18
Step function	1.94 - 2.24	1.29 - 1.46	1.20 - 1.36	1.09 - 1.23

80s stand out as particularly costly, but now we also see a pronounced spike in 2001, driven by US deaths in the 9/11 attacks, and then a rise later in the 2000s due to the US wars with Afghanistan and Iraq. The dark and light green lines show the total annual cost of conflict where the multipliers are instead imputed according to the step function. By design, this imputation is much more conservative and it is easy to see that the green lines are well above the others. Finally, the blue line shows the global annual total of military spending. There is a clear trend upward through the Cold War and a fall-off after its end, followed by a renewed increase in the early 2000s. Remarkably, in almost every year in the data, the cost from arming is much greater the cost from conflict that year. The sole exceptions are in the late 60s and, just barely, in 2001 according to the more conservative step-function imputation of multipliers.

Considering the whole half-century of arming and conflict in aggregate, Table 2 reports the ratio of arming cost to conflict cost under each method of imputing multipliers and across different assumptions about war-induced arming. The range in each cell represents the uncertainty deriving from how to partition battle-deaths across belligerents. Remarkably, the cost of all arming is larger than the cost of all conflicts even under the most conservative method and assumptions. If we exclude war-induced arming from the cost of war, the cost of arming outweighs the cost of war by at least a factor of two or three.

Our conclusion that the global cost of arming exceeded that of conflict over the last half-century (up to 2008) has two potential failure points that seem obvious to us. The first is

the assumption that the battle-deaths-to-total-cost multiplier should be tightly coupled to a belligerent's level of development, so that we can use the latter to impute values for the former. We offered a theoretical argument for this assumption in the previous section, and showed that it held for the capital-to-labor ratio exhibited in states' militaries, which should theoretically be closely related to the multiplier. The results in Table 2 show that even a quite relaxed version of this assumption is sufficient to support the conclusion that arming costs more than war. Recall that the step-function method of imputing multipliers requires only that the multiplier tends to rise in the level of development—for every unknown case, it imputes the highest possible multiplier that is consistent with that rise. Because, even under this extraordinarily conservative method, arming still costs more than war, we are reassured about this potential point of failure for our conclusion.

A second obvious point of failure is in our reference multipliers: the multipliers we calculated from the four cases in which we have comprehensive estimates of the total cost of particular conflicts for particular belligerents. The authors of those estimates freely admit to considerable uncertainty in their accounting. What if their estimates are too low? That would lead to too-low multipliers, which might in turn cause our imputed multipliers for other cases to be too low. To address this, we conducted a sensitivity analysis on our conclusion. How far off does each reference case's cost estimate have to be to reverse our conclusion, either individually or collectively?

Table 3 shows by what factor each reference case's multiplier would have to be too-low to upset our finding that arming has cost more than conflict, under each method of imputing multipliers. Here, we exclude all wartime arming from the cost of war, taking it to instead be part of the cost of arming. The table makes clear that there is room for either or both reference multipliers to be off by a very large margin without altering the finding that arming costs more than war. Of course, if either reference case's cost estimate is actually too high, correcting this would only strengthen our finding.

Table 3: Sensitivity to Underestimates of Reference Multipliers

Method of Multiplier Imputation	multiplier(s) would have to be too low by a factor of:		
	Iraq	US	Both
Linear	3.5	4.0	2.5
Step function	3.5	7.5	4.0

Implications

Our finding that the global costs of arming exceed those of war yields important implications for both scholars and policymakers. For scholars, it suggests a new research agenda intended to strengthen (or refute) the evidence for this finding and to investigate the possible policy instruments for reducing arming. For policymakers, it reveals the potentially large gains humanity could accrue from a new attempt at multilateral arms control, even relative to further efforts to constrain or prevent war.

We acknowledge that our finding must be regarded as tentative. It depends on a theoretical assumption—that the ratio of battle-deaths to other resources a state loses in conflict tends to fall in its level of development—and a very small number of cases for which comprehensive estimates of the cost of war are available. We have endeavored to show that this assumption is theoretically defensible and borne out in states' military capital-to-labor ratios, and also to demonstrate that our finding is robust even to large errors in the case estimates. But there is no getting around the fact that the theoretical assumption is just that, and we cannot rule out the possibility that the case estimates are so far off the mark that correct estimates would lead to the opposite finding.

Even so, the potential importance of this finding should motivate new research aimed at reinforcing or overturning it. The easiest way to do this would be to develop comprehensive estimates of the cost of war for additional cases. Every additional case for which the total cost of war is plausibly estimated provides an additional test of the assumption that the lives-

to-all-resources-lost multiplier rises in development, and gives us more information about the shape of the relationship between them. Each new case also reduces the dependence of our calculation on getting the other comprehensively-estimated cases right, lessening concern about potentially large errors in those estimates. The most valuable cases to analyze would be those in which the belligerents were less developed than Iran, Iraq, and the United States, because most conflicts involve such states and we currently have no estimates for them. These cases would be especially informative if the belligerents also did not possess unusually advanced military technology relative to their level of development. This would reduce the bias in our calculation that comes from relying on cases that are outliers in this respect.

We do not mean to slight the difficulty of developing additional comprehensive estimates of the cost of particular wars, especially among poorer belligerents, which tend to have fewer administrative records on which to base such estimates. But we do believe such estimates can be constructed, and we suspect that they have not been made only because scholars did not see why doing so would be worth the effort.

We also need more thorough measures of the cost of arming. Existing data on military spending are valuable, but they ignore the many states that rely at least partly on conscription to acquire resources for their militaries, and the cost this imposes on their economies. It is important to develop measures of the degree of conscription across states and years, and to do so for both labor (as conscription is colloquially understood) and capital. Government commandeering of physical plant and involuntary conversion of factories to arms production is an important cost of arming in many instances. As it stands, measures of the cost of arming based solely on military spending inevitably underestimate the true cost of arming by an unknown amount that probably varies across states and over time.

Recognizing this, it is important to reiterate that our measures of both the cost of war and the cost of arming are surely biased, and probably by large margins: by design, the true

cost of war is likely substantially lower than the figures we compute while the true cost of arming is likely much higher. Such biased estimates are good enough for the purpose of this paper, which was to show that the biased-upward measure of the cost of war is smaller than the biased-downward measure of the cost of arming. However, they should not be used by subsequent scholars as unproblematic estimates of the cost of either.

Assuming our finding survives new evidence and additional testing, it raises the value of investigating the determinants of arming and especially the possibilities for reducing it. Why does the level of arming vary from country to country and year to year? What leads to an acceleration or deceleration of arming, either by individual states or collectively? What are the conditions under which multilateral agreements to limit arming can be concluded and enforced? Why do some eras witness extensive efforts at arms control while others see only neglect? There is some research already on each of these questions, but none has received the kind of sustained attention from scholars that has been accorded to civil and interstate conflict.

Most importantly, our finding suggests that the time may be ripe for a new push for comprehensive arms control. This was last attempted at the beginning of the Cold War between the superpowers, and with somewhat greater success by the major powers during the interwar period. In the last thirty years, arms control has mostly been narrowly focused on weapons of mass destruction—a worthy and arguably successful effort, but one that does little to reduce the cost of arming given the modest cost of those particular weapons. The scope for welfare gains from a new set of agreements focused on conventional forces, which make up the bulk of arms spending around the world, is tremendous, and our study suggests it may be even greater than the potential gains from further measures to reduce the occurrence or severity of war.

In principle, it should always be possible to implement multilateral reductions in arms that leave the distribution of power relatively undisturbed but save the participants the

money that would otherwise pay for those arms. In practice, monitoring agreed reductions in a way that assures compliance but does not otherwise threaten participants' security is sometimes challenging (Coe and Vaynman, 2020). But potential monitoring arrangements for comprehensive controls on conventional forces have simply not been investigated in recent decades, with most analyses focused on WMD or new weapons in cyberspace or orbit. It is not clear one way or the other whether recent advances in espionage and sensing arising from the revolution in information and communication technology would make such monitoring easier than it was in the past.

Even if a renaissance in multilateral arms control were only partially successful, the resulting gains in human welfare would be immense. The world currently spends more than two cents of every dollar on its militaries. Reducing that by half would free up enough money to end world poverty, provide bed nets to every person exposed to mosquito-borne disease, treat every person infected with HIV, and increase the world's investment in research and development by half (Coe and Vaynman, 2020). Every year the world neglects reductions in arms, $2.7 trillion dollars are lost to these and many other worthier causes.

References

Allansson, Marie, Erik Melander and Lotta Themnér. 2017. "Organized Violence, 1989–2016." *Journal of Peace Research* 54(4):574–587.

Brooks, Stephen G. and William C. Wohlforth. 2015. "The Rise and Fall of the Great Powers in the Twenty-First Century: China's Rise and the Fate of America's Global Position." *International Security* 40(3):7–53.

Coe, Andrew J. 2011. "Costly Peace: A New Rationalist Explanation for War." Unpublished manuscript.

Coe, Andrew J. and Jane Vaynman. 2020. "Why Arms Control Is So Rare." *American Political Science Review* 114(2):342–355.

Collier, Paul. 1999. "On the economic consequences of civil war." *Oxford economic papers* 51(1):168–183.

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

Fearon, James D. 1995. "Rationalist Explanations for War." *International Organization* 49(03):379–414.

Fearon, James D. 2018. "Cooperation, Conflict, and the Costs of Anarchy." *International Organization* 72(3):523–559.

Gates, Scott, Håvard Hegre, Håvard Mokleiv Nygård and Håvard Strand. 2012. "Development consequences of armed conflict." *World Development* 40(9):1713–1722.

Ghobarah, Hazem Adam, Paul Huth and Bruce Russett. 2003. "Civil Wars Kill and Maim

People—Long After the Shooting Stops." *American Political Science Review* 97(2):189–202.

Ghobarah, Hazem Adam, Paul Huth and Bruce Russett. 2004. "The Post-war Public Health Effects of Civil Conflict." *Social science & medicine* 59(4):869–884.

Gleditsch, Nils Petter, Peter Wallensteen, Mikael Eriksson, Margareta Sollenberg and Håvard Strand. 2002. "Armed Conflict 1946–2001: A New Dataset." *Journal of Peace research* 39(5):615–637.

Goldstein, Joshua S. 2011. *Winning the War on War: The Decline of Armed Conflict Worldwide*. Dutton.

Hathaway, Oona A. and Scott J. Shapiro. 2017. *The Internationalists: How a Radical Plan to Outlaw War Remade the World*. Simon & Schuster.

Honaker, James and Gary King. 2010. "What to Do about Missing Values in Time-Series Cross-Section Data." *American Journal of Political Science* 54(2):561–581.

Kugler, Tadeusz. 2017. Demography and International Relations: Economics, Politics, Sociology, and Conflict. In *Advancing Interdisciplinary Approaches to International Relations*. Springer chapter 8, pp. 229–261.

Lacina, Bethany and Nils Petter Gleditsch. 2005. "Monitoring Trends in Global Combat: A New Dataset of Battle Deaths." *European Journal of Population/Revue européenne de Démographie* 21(2):145–166.

Martí Sempere, Carlos. 2018. "What Is Known about Defence Research and Development Spill-Overs?" *Defence and Peace Economics* 29(3):225–246.

Meirowitz, Adam and Anne E. Sartori. 2008. "Strategic Uncertainty As a Cause of War." *Quarterly Journal of Political Science* 3(4):327–352.

Mofid, Kamran. 1990. *The Economic Consequences of the Gulf War*. Routledge.

Moretti, Enrico, Claudia Steinwender and John Van Reenen. 2019. The Intellectual Spoils of War? Defense R&D, Productivity and International Spillovers. Technical report National Bureau of Economic Research.

Office of Management and Budget. 2019. "Historical Tables." Accessed at https://www.whitehouse.gov/omb/historical-tables/.

Organisation for Economic Co-operation and Development. 2022. "Main Science and Technology Indicators.".

Powell, Robert. 1993. "Guns, Butter, and Anarchy." *American Political Science Review* 87(1):115–132.

Singer, J. David. 1987. "Reconstructing the Correlates of War Dataset on Material Capabilities of States, 1816–1985." *International Interactions* 14:115–132.

Singer, J. David, Stuart Bremer and John Stuckey. 1972. Capability Distribution, Uncertainty, and Major Power War, 1820–1965. In *Peace, War, and Numbers*, ed. Bruce Russett. Sage pp. 19–48.

Slantchev, Branislav L. 2005. "Military Coercion in Interstate Crises." *American Political Science Review* 99(4):533–547.

Stiglitz, Joseph E. and Linda J. Bilmes. 2008. *The Three Trillion Dollar War: The True Cost of the Iraq Conflict*. W. W. Norton & Company.

Vaynman, Jane. 2022. "Enemies in Agreement: Domestic Politics, Uncertainty, and Cooperation between Adversaries." Working manuscript.

Viscusi, W. Kyp. 2018. *Pricing Lives: Guideposts for a Safer Society*. Princeton University Press.