School of Arts and Social Sciences
CITY UNIVERSITY LONDON

Department of Economics

Allocating Security Expenditures under Knightian Uncertainty: an Info-Gap Approach

Michael Ben-Gad[1]
City University London

Yakov Ben-Haim
Faculty of Mechanical Engineering, Technion-Israel Institute of Technology, Haifa, Israel

Dan Peled
Department of Economics, University of Haifa, Israel

Department of Economics
Discussion Paper Series
No. 15/09

[1] Corresponding author: Michael Ben-Gad, Department of Economics, City University London, Northampton Square, London EC1V 0HB, UK. Email: Michael.Ben-Gad.1@city.ac.uk

Allocating Security Expenditures under Knightian Uncertainty: an Info-Gap Approach*

Michael Ben-Gad[a] Yakov Ben-Haim[b]

Dan Peled[c]

December 2015

Abstract

We apply the information gap approach to resource allocation under Knightian (non-probabilistic) uncertainty in order to study how best to allocate public resources between competing defense measures. We demonstrate that when determining the level and composition of defense spending in an environment of extreme uncertainty *vis-a-vis* the likelihood of armed conflict and its outcomes, robust-satisficing expected utility will usually be preferable to expected utility maximization. Moreover, our analysis suggests that in environments with unreliable information about threats to national security and their consequences, a desire for robustness to model misspecification in the decision making process will imply greater expenditure on certain types of defense measures at the expense of others. Our results also provide a positivist explanation of how governments seem to allocate security expenditures in practice.

JEL classification: D81; F51; F52

Keywords: Defense; Knightian Uncertainty; Robustness; Info-gap

[a]Department of Economics, City University London, Northampton Square, London EC1V 0HB, UK.

[b]Faculty of Mechanical Engineering, Technion-Israel Institute of Technology, Haifa 32000, Israel.

[c]Department of Economics, University of Haifa, Haifa 31905, Israel.

*This research was supported by the Economics of National Security Program, Samuel Neaman Institute, grant no. 358. Ben-Gad: *mbengad@city.ac.uk*, Ben-Haim: *yakov@technion.ac.il*, Peled: *dpeled@econ.haifa.ac.il*

"*There are things we know that we know. There are known unknowns. That is to say there are things that we now know we don't know. But there are also unknown unknowns. There are things we do not know we don't know. So when we do the best we can and we pull all this information together, and we then say well that's basically what we see as the situation, that is really only the known knowns and the known unknowns. And each year, we discover a few more of those unknown unknowns.*[1]

Donald Rumsfeld
Former U.S. Secretary of Defense

1 Introduction

What is the best procedure for a policy maker to employ when allocating scarce resources to both deter military aggression, and in the event deterrence is unsuccessful, mitigate its effects, if previous experience provides little relevant guidance? This paper demonstrates how both the scope and nature of defense expenditure should change, if policy makers wish their decisions to be robust to the type of Knightian (non-probablistic) uncertainty they ordinarily confront.

When allocating resources, policy makers, like households, face trade-offs. Many of the trade-offs households encounter can reasonably be modeled in a deterministic setting—their incomes may be uncertain, but when deciding between a new refrigerator or a new television set, we generally assume that consumers know the marginal utility they will derive from each. By contrast, policy makers often confront decisions in which the connection between allocations, and the desirability of outcomes, are uncertain. Nonetheless they are often able to derive reliable probabilistic models that link different allocations of funding with the moments of a distribution of outcomes. For example, no one can predict precisely by how much a dollar transferred between different components of the public health budget will affect an individual citizen's longevity. However, policy makers know a great deal about not only the prevalence, but the distribution of infectious diseases, heart problems, and cancer for various population groups and the efficacy of different treatments for large samples of patients.

Unlike public health decisions, allocating resources for national security involves decisions where experiments are not possible (or at least unwise). Moreover, as relations between any set of political actors on the international stage are constantly shifting, and both military doctrines and technology continue to evolve, previous experience may not provide much useful information upon which to base present-day decisions on how best to cope with future threats.[2] The policy

[1] Antulio J. Echevaria II (2008).

[2] David Hume, writing in 1748: "What is the foundation of all conclusions from experience?.....All experimental

maker whose country is threatened by hostile forces may know little about the probability distribution of possible damage or losses his country could suffer in the event of armed conflict, as any estimate will require both an intimate knowledge of the adversary's capabilities and tactics, and the degree to which they might be counteracted by different defense allocations. Indeed, the probability that an adversary will launch an attack may itself be unknown to the policy maker, as it requires judgements about the intentions and expectations of foreign leaders, or sub-state actors and yet is still dependent on the type of defense capabilities our policy maker has chosen. Finally, often enough the enemy's own political and military decision makers are uncertain about their own goals or the best means to achieve them, at the time the our decision maker must determine the best allocation.

When allocating resources for national defense, the maximization of expected utility will often prove inadequate because it assumes that probability distributions are known, or at least that statistical moments can be calculated. Required is a methodology that enables the policy maker to allocate resources without requiring the use of unavailable probabilistic information. By implementing the info-gap methodology, policy makers combine what they know about security threats, and their relationships to installed military capabilities with specified policy requirements, without requiring the policy maker to know how wrong the available information is. Moreover, our analysis highlights one rationale for heightened spending on some types of defense measures when there is little reliable information about the random nature of threats to national security. This provides a possible positivist explanation for the way governments allocate these expenditures today, and how they have done so in the past.

In section 2 we describe the general formulation of the info-gap methodology in the context of defense expenditure. In section 3, by way of example, we consider a simple version of a dilemma faced by many policy makers today: how much to spend overall on defense and how much of that spending should be diverted towards the precision munitions, electronic warfare, and space-based communications and intelligence necessary to implement the 'Revolution in Military Affairs' (RMA) doctrine, and away from continued investment in large units employing traditional industrial-age military technology and platforms.

---

conclusions proceed upon the supposition that the future will be conformable to the past....why this experience should be extended to future times, and to other objects, which for aught we know, may be only in appearance similar; this is the main question on which I insist....If there be any suspicion that the course of nature may change, and that the past may be no rule for the future, all experience becomes useless, and can give rise to no inference or conclusion."

[2]First elaborated during the 1980's by Soviet military theorists, in particular Marshal Nikolai Ogarkov, then chief of the Soviet General Staff, the RMA doctrine of war has already had a profound effect on Western military planning. The move to adapt the U.S. military to RMA type warfare was spearheaded by the Department of Defense's Office of Net Assessment, and the man who has headed it since 1973, Andrew Marshall, and the vice chairman of the Joint Chiefs of Staff between 1994 to 1996, Admiral William Owens. See Eliot A. Cohen (1996).

Throughout this paper, the risk of being attacked and the probability distribution of damage conditioned on an attack occurring—both of which are info-gap uncertain—are related to both the size and structure of the military force policy makers have chosen. We demonstrate that when policy makers must operate in an environment in which information about the relationship between threats and different allocations for defense is unreliable, expected utility maximization does not provide the best guidance for making decisions, because it is completely non-robust to model misspecification. Using a hypothetical specification of security risks and the cost effectiveness of alternative defense measures, we show that the desire of policy makers to attain some degree of robustness to model misspecification, will have profound effects on both the level of resources devoted to defense, and the distribution of these expenditures across the two possible types of defense measures.

2 Formulation

Consider a country facing various threats to its national security. These threats can take various forms, including invasion by an aggressive neighbor or a terrorist attack. Policy makers perceive threats to security from all sources as a bivariate distribution that includes both the event of being attacked, and the damage the country will sustain conditional on the attack taking place. In contrast to standard stochastic optimization problems, the distribution itself is uncertain. Policy makers must decide what portion of the economy's resources they will devote to countering security risks, as well as how to allocate this expenditure between different defense measures, where each measure has a different effect on both the risk of an attack and the potential damage it will inflict. In this paper, this defense resource allocation dilemma is embedded within a standard economic framework in which the representative risk-averse individual in this country derives utility $u(c)$ only from consumption, c. The threats considered here, and the effect of any countermeasures, affect the resulting level of utility through their influence on the resources that remain for consumption, the likelihood of being attacked, and the conditional distribution of losses an attack may inflict.

Normalizing the economy's resources to 1, the policy maker must choose the fraction of all resources to devote to each of N different risk-mitigating expenditures $\chi = (\chi_1, \ldots, \chi_N)$. Without government debt, we require $\sum_i^N \chi_i \leq 1$. Any government expenditure detracts from the resources available for consumption, so $c = 1 - \sum_i^N \chi_i$. On the other hand these government expenditures reduce the likelihood of an attack and the fractional loss ψ in resources resulting from an attack if it takes place, where $0 \leq \psi \leq 1$.

The probability density function (pdf) of realized threats, conditioned on risk mitigating expenditures, is $p(\psi|\chi)$—a probability distribution unknown to policy makers. The best available estimate of $p(\psi|\chi)$ is denoted $\widetilde{p}(\psi|\chi)$, but it is incontrovertible that $\widetilde{p}(\psi|\chi)$ is highly unreliable. We denote the probability of attack, as a function of defense spending χ, by $P_{\rm w}(\chi)$, and the best (but highly uncertain) estimate of this function is $\widetilde{P}_{\rm w}(\chi)$. The salient feature of this function, (illustrated in Figure 2 for N=2), is that $\widetilde{P}_{\rm w}(\chi)$ is a decreasing function of each type of defense expenditure, holding expenditures on other types fixed.

Let $R(\chi|p, P_{\rm w})$ be the expected utility resulting from defense expenditure χ, when the probability of the attack being realized is $P_{\rm w}$, and the pdf of the damage ψ is $p(\psi|\chi)$:

$$R(\chi|p, P_{\rm w}) \quad = \quad P_{\rm w}(\chi) \int_0^1 u \left[(1-\psi)(1-\sum_i^N \chi_i) \right] p(\psi|\chi)\, {\rm d}\psi + (1-P_{\rm w}(\chi)) u(1-\sum_i^N \chi_i) \quad (1)$$

where the last term, $u(1-\sum_i^N \chi_i)$, represents the level of utility agents in this economy enjoy if no security risks materialize.

The info-gap model in the current context is an unbounded family of nested sets of probability models of $p(\psi|\chi)$ and $P_{\rm w}(\chi)$, indexed by α, which represents the unknown horizon of uncertainty in the policy maker's best estimate of the chances of an attack occurring and of the conditional distribution of the damages it would inflict. We denote this info-gap model by $\mathcal{F}[\alpha, \widetilde{p}(\psi|\chi), \widetilde{P}_{\rm w}(\chi)]$, where $\alpha \geq 0$.[3] The important feature of this info-gap model is that larger α entails larger possible deviations of the actual distributions of threats from their best estimates, and accordingly a wider range of actual challenges the defense allocation chosen will need to confront. A specific illustration of the info-gap model will be presented in section 3.2.

The policy maker requires that the expected utility be no less than a critical value, R_c. This critical value may be uncertain, or the policy maker may wish to identify an acceptable value of R_c that can be confidently anticipated, given a proposed allocation χ. In other words, the goal of the policy maker is to maximise robustness to model uncertainty subject to a minimally acceptable expected utility. That is, the policy makers wishes to choose an allocation decision χ that yields the widest margin of error in model specification under which the expected utility equals or exceeds R_c.

Specifically, the robustness of allocation χ is the maximum horizon of uncertainty, α, up to

[3]Info-gap models obey two axioms:

(i) *Nesting* asserts that the range of possible pdfs increases as α increases:

$$\alpha < \alpha' \quad \Longrightarrow \quad \mathcal{F}[\alpha, \widetilde{p}, \widetilde{P}_{\rm w}] \subseteq \mathcal{F}[\alpha', \widetilde{p}, \widetilde{P}_{\rm w}] \quad (2)$$

(ii) *Contraction* asserts that when $\alpha = 0$, the best estimated models are the only possibilities:

$$\mathcal{F}[0, \widetilde{p}, \widetilde{P}_{\rm w}] = \left\{ \widetilde{p}, \widetilde{P}_{\rm w} \right\} \quad (3)$$

These two axioms endow α with its meaning of a *horizon of uncertainty.*

which the expected utility is no less than R_c for all probability functions in $\mathcal{F}(\alpha)$:

$$\widehat{\alpha}(R_c, \chi) = \max\left\{\alpha \mid \; R(\chi|p, P_w) \geq R_c, \; \forall \; (p, P_w) \in \mathcal{F}(\alpha)\right\} \tag{4}$$

More robustness is better than less, so the robust-optimal allocation at the minimally acceptable reward R_c is the allocation that maximizes the robustness:

$$\widehat{\chi}(R_c) = \arg\max_{\chi} \widehat{\alpha}(R_{\mathrm{c}}, \chi) \tag{5}$$

In summary, $\widehat{\alpha}(R_c, \chi)$ is the robustness of allocation χ for achieving expected utility no less than R_c. Likewise, $\widehat{\alpha}^*(R_c)$ is the maximal robustness for achieving expected utility no less than R_c by choosing the appropriate values of χ. It is the maximum value of $\widehat{\alpha}(R_c, \chi)$ from (4), evaluated at R_c and $\widehat{\chi}(R_c)$ from (5).

$$\widehat{\alpha}^*(R_c) = \max_{\chi}\{\alpha \mid R(\chi|p, P_{\mathrm{w}}) \geq R_c, \quad \forall \quad (p, P_{\mathrm{w}}) \in \mathcal{F}(\alpha, \widetilde{p}, \widetilde{P}_{\mathrm{w}})\} \tag{6}$$

As in any info-gap model (Ben-Haim, 2006), there is a fundamental trade-off between minimally acceptable reward and robustness to uncertainty: the former decreases as the latter increases as will be demonstrated subsequently.

3 An Illustration with Two Types of Security Expenditure

3.1 Basic Structure

By way of example, suppose all military expenditures fall into one of two broad categories. First, there is the expenditure that incorporates recent innovations in military technology and tactics, based on the intensive use of information technology, high-precision weaponry, and satellites employed for both intelligence and command and control systems. We denote this type of defense expenditure, associated with the 'Revolution in Military Affairs' (RMA) doctrine, as χ_1. Second is the development and maintenance of large military formations, composed of large numbers of troops receiving traditional military training to serve as infantrymen or to operate armor, artillery, battleships and bombers. We denote this, the more traditional, industrial-age type of military expenditure, and the one most closely associated with twentieth-century warfare, as χ_2. Both χ_1 and χ_2 are measured as shares of GDP, and we assume there is no possibility of international borrowing. Recall that $p(\psi|\chi)$ is the probability density function (pdf) of the damage the nation sustains in terms of lost GDP in the event that it is attacked, and assume its shape is influenced by both the overall quantity of resources devoted to security, $\chi_1 + \chi_2$, and also by how resources are divided between χ_1 and χ_2. Next we illustrate the kinds of considerations which can be employed to set the shape of the probability models used in the sequel.

With its reliance on small, highly mobile and specially trained military units, the RMA doctrine can under the most optimal conditions limit the severity of casualties and other losses for many kinds of attacks, but is more prone to catastrophic failure, particularly in the event of a massive invasion by a large military force. Accordingly, we assume that holding the GDP share of total defense expenditures fixed, the more defense expenditures are weighted towards RMA type, at the expense of traditional weapons systems and military formations, the lower are the expected losses conditioned on an attack occurring, but the risk of extremely high damages increases. By contrast, traditional large military formations can insure against the most severe and widespread losses, but reliance on them will entail higher losses under most circumstances. With all defense expenditure restricted to fall into these two broad categories, policy makers' best (but highly uncertain) estimate of the damage pdf, conditional on being attacked and given allocation of defense expenditure χ_1 and χ_2, is $\widetilde{p}(\psi|\chi_1, \chi_2)$. Given the long planning horizon necessary to prepare defense forces, the relevant unit of time in this model is a decade.

The functional form for the best estimate of the damage density function reflects all available knowledge about possible losses incurred in the event of suffering an attack, and how the risks of these losses relate to the values of χ_1 and χ_2. In the sequel we provide the specification of the particular probability density function we chose for our illustration. Here, it suffices to note the two most salient features implied by our specification, which capture the essence of the tradeoff assumed to exist between the two alternative types of defense expenditure:

1. Mean damage generally declines in each of the security expenditure types. That is, at least in the neighborhood of $\widehat{\chi}(R_c)$, we assume:
$$\frac{\partial \mathrm{E}(\psi)}{\partial \chi_i} < 0,\ i \in \{1, 2\} \tag{7}$$
2. At the same time, while holding total security expenditure fixed, we expect to observe that increases in the share devoted to traditional security expenditures lowers the probability of extreme damage, defined here as 60% of output. That is, we expect to generally observe:
$$\left.\frac{\partial \mathrm{Prob}(\psi < 0.6)}{\partial \chi_2}\right|_{\chi_1+\chi_2} > 0 \tag{8}$$

In Figure 1 we graphically present the density function $\widetilde{p}(\psi|\chi)$, based on the specification in the appendix. In that figure, total defense expenditures $\chi_1 + \chi_2$ is set at a constant ten percent of GDP, while the value of the traditional military expenditure χ_2 varies between the values 0 and 0.1.

These damage estimates are conditional on the occurrence of an attack, but what do we know about its probability?[4] Clearly, determining the probability of attack will necessitate estimating the military capabilities of potential adversaries—capabilities they may varyingly wish to conceal or exaggerate.[5] More difficult still is the need to assess the intentions of foreign leaders, whose decision processes and perceptions are often informed by cultures and historical memories different from one's own.

> The assumption that all actors in the arena of world politics can be counted upon to take a rational approach in foreign policy is not without some validity and utility; but it does not offer a sufficient basis for estimating how these other actors view events, calculate their options, and make their choices of action....To describe behavior as "rational" is to say little more than that the actor attempts to choose a course of action that he hopes or expects to further his values. But, of course, what the opponent's values are and how they will affect his policymaking and decisions in different kinds of situations remains to be established. Moreover, foreign-policy issues are typically complex in that they raise multiple values and interests that cannot easily be reconciled. Even for the rational actor, therefore, choice is often difficult because he faces a value trade-off problem. How the opponent will resolve that dilemma is not easily foreseen, even by those close to him let alone by those in another country who are attempting to predict his action on the basis of the rationality of the assumption.[6]

Yet even if it were possible to understand a potential opponent's existing intentions in a manner sufficient to assign probabilities to an attack as a function of one's own defense allocations, this would still be insufficient. Intentions change.

Though the Japanese were not privy to German plans to attack the Soviet Union, during the months before Operation Barborossa was launched on 22 June, 1941, the Japanese actively considered beginning to prepare for a 'northern war', an attack on the Soviet Union's far eastern territories. Throughout that summer, civilians and military people in the Japanese government weighed the various merits of attacking the Soviet Union, attacking the Western powers to the

[4] "As the ancient retiree from the Research Department of the British Foreign Office reputedly said, after serving from 1903-50: 'Year after year the worriers and fretters would come to me with awful predictions of the outbreak of war. I denied it each time. I was only wrong twice.'" Thomas L. Hughes, 1976, p. 48.

[5] "An adversary can still decide to attack even though his capabilities are relatively weak (1) if he miscalculates the strength of the intended victim (as did the Germans in their attack on the Soviet Union in 1941, or the Arab States in their underestimation of Israeli capabilities in 1967); (2) if he is more interested in applying political pressure or making political gains even at the cost of military defeat; (3) if he gambles that his surprise attack will have a force multiplier effect sufficient to compensate for his inferior capabilities." Handel, 1989, p. 241.

[6] Alexander George, 1980, p. 66.

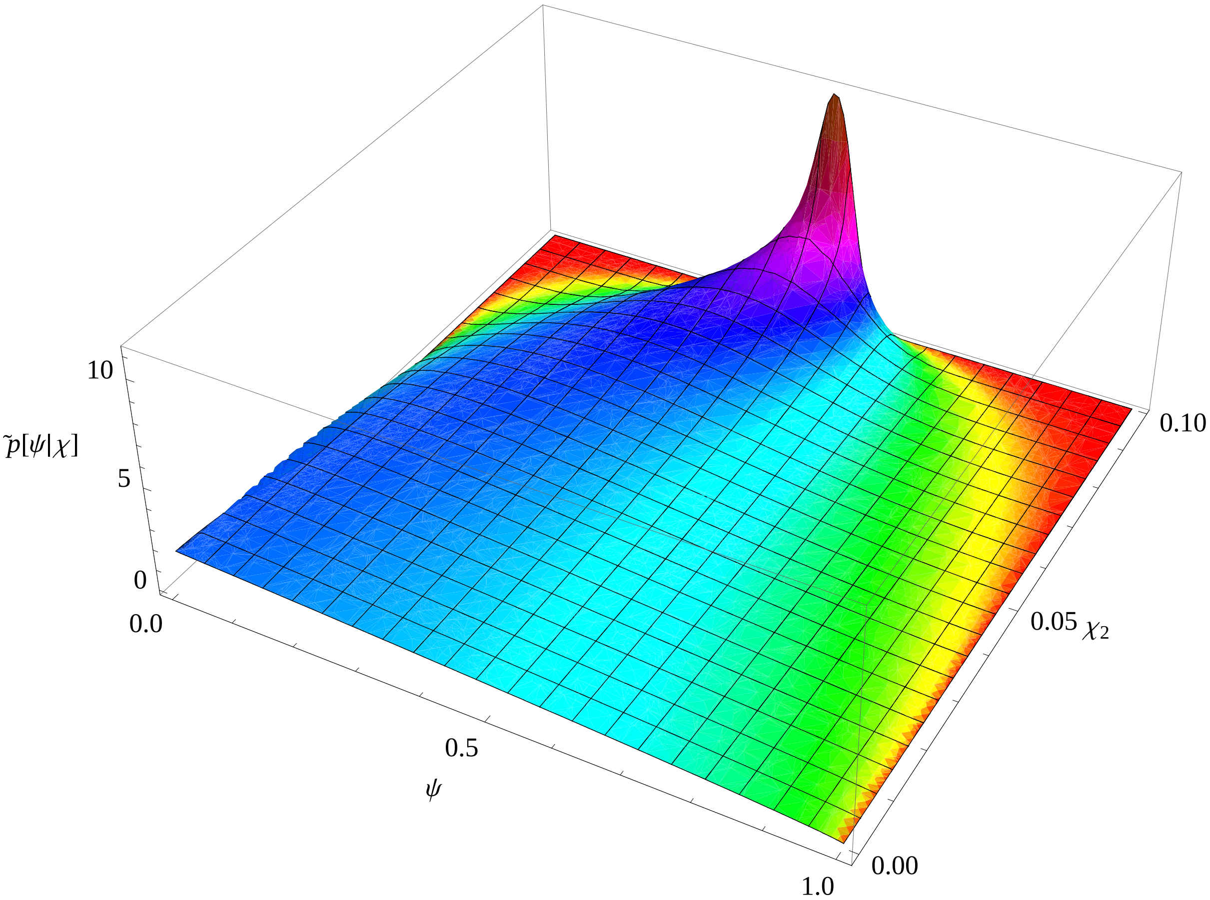

Figure 1: The best estimate of the probability density of damage, conditional on defense expenditure $\widetilde{p}(\psi|\chi)$ as in (26), for different values of χ_2, traditional military expenditure, where total defense spending comprising both RMA-type and traditional expenditure is $\chi_1 + \chi_2 = 0.1$.

south, or continuing to rely on diplomacy while consolidating gains in China and newly acquired French Indo-China. Only at a military conference on 3 September did the more aggressive view of the Japanese Army prevail over that of the more cautious Navy to initiate a war in the south, in response to continuing United States refusal to restore normal economic relations and not impede further Japanese ambitions in China. Three days later an imperial council approved the military's recommendation that "if, by the early part of October, there is still no prospect of being able to attain our demands, we shall immediately decide to open hostilities against the United States, Great Britain, and the Netherlands."[7] The decision to begin preparations for the 'southern war' commenced only three months and a day prior to the attack on 7 December, 1941. Hence the probability of an attack over the course of a decade is subject to great uncertainty—policy makers can scarcely quantify with much confidence the probability of an attack long before an enemy might even be considering one. Figure 2 portrays the potential subtle influences defense expenditures of both types can have on the probability that an attack will be launched.

3.2 A Fractional Info-Gap Model of Uncertainty

The density function $\widetilde{p}(\psi|\chi)$ is the best estimate of the pdf of damage from an attack, given security allocation $\chi = (\chi_1,\, \chi_2)$. However, this estimate is based on fragmentary and unreliable evidence, and hence contains potentially serious but unidentifiable errors. The same is true for the estimated probability of attack, $\widetilde{P}_{\rm w}(\chi)$. The true values deviate from these estimates by unknown amounts. We use a *fractional error* info-gap model to represent the gaps in both the pdf of the damage and the probability of attack (Ben-Haim, 2006). Let $\mathcal{P}$ denote the set of all pdfs on [0, 1]. For any $\alpha \geq 0$ our info-gap model consists of all density functions and probability values that differ proportionally from $\widetilde{p}(\psi|\chi)$ and $\widetilde{P}_{\rm w}(\chi)$, respectively, by no more than α. That is:

$$\mathcal{F}(\alpha, \widetilde{p}, \widetilde{P}_{\rm w}) = \Big\{ \, (p(\psi),\, P_{\rm w}) : \quad \begin{array}{l} p(\psi) \in \mathcal{P},\ |p(\psi) - \widetilde{p}(\psi|\chi)| \leq \alpha \widetilde{p}(\psi|\chi),\ \text{for all } \psi \\ 0 \leq P_{\rm w} \leq 1,\ |P_{\rm w} - \widetilde{P}_{\rm w}(\chi)| \leq \alpha \widetilde{P}_{\rm w}(\chi), \alpha \geq 0 \end{array} \Big\}, \tag{9}$$

The value of the fractional error α is unknown, so the info-gap model is not a single set, but rather an unbounded family of nested sets of possible pdfs and probabilities. Since α is unbounded from above, there is no known worst case.

3.3 Robustness Function

Conditioned on an attack occurring, the expected utility under the best estimates of damage density is:

[7] Butow, 1961, p. 250.

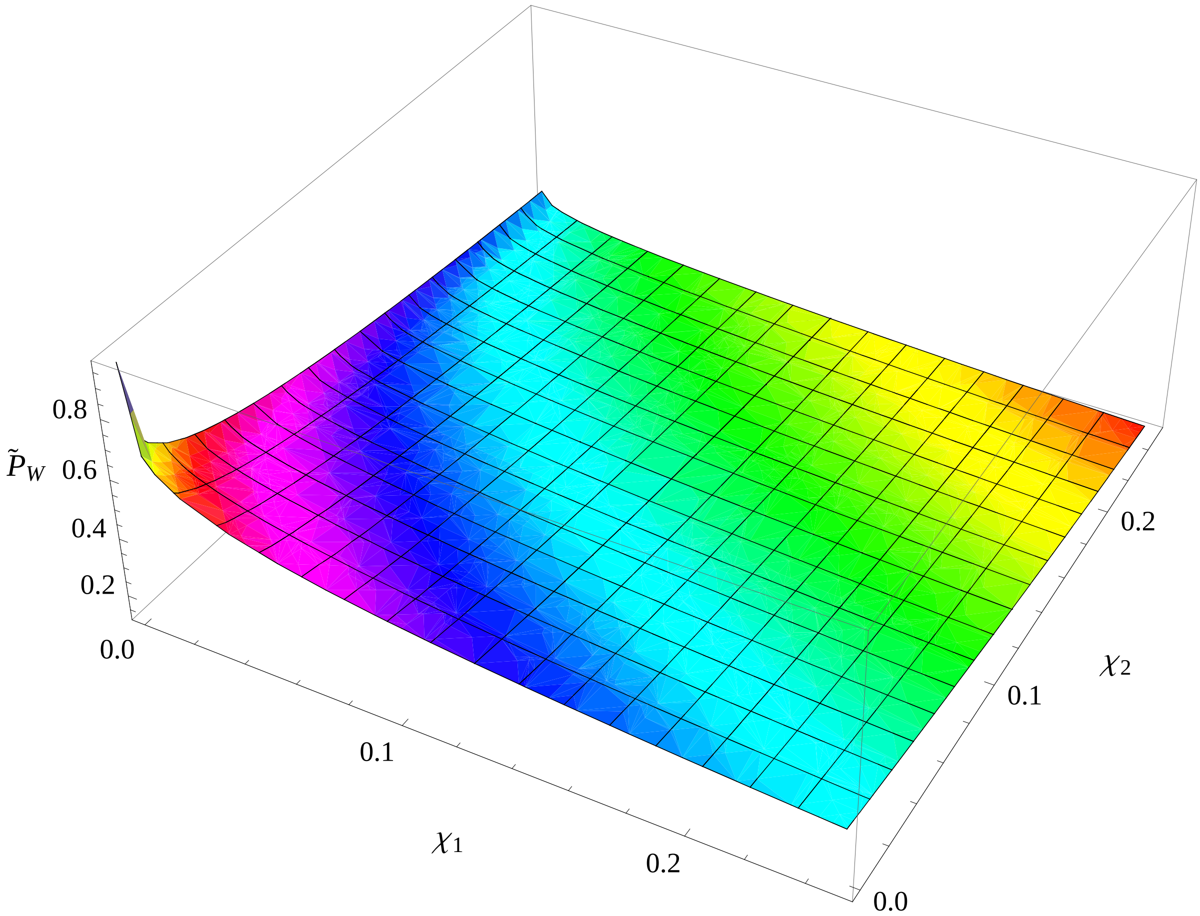

Figure 2: The best estimate of the probability of being attacked $\widetilde{P}_{\mathrm{w}}(\chi)$ (as in (30)) for different values of χ_1 and χ_2.

$$\widetilde{r}(\chi) = \int_0^1 u\left[(1-\psi)(1-\chi_1-\chi_2)\right]\widetilde{p}(\psi|\chi)\,\mathrm{d}\psi \tag{10}$$

The unconditional expected utility, based on the best estimates of the damage pdf $\widetilde{p}(\psi|\chi)$ and the probability of attack $\widetilde{P}_{\mathrm{w}}(\chi)$, (as defined in (1)) is:

$$\widetilde{R}(\chi) = \widetilde{P}_{\mathrm{w}}(\chi)\widetilde{r}(\chi) + (1-\widetilde{P}_{\mathrm{w}}(\chi))u(1-\chi_1-\chi_2) \tag{11}$$

We assume that $\widetilde{r}(\chi) < u(1-\chi_1-\chi_2)$, so that the expected utility in the event of an attack is always smaller than it is in the absence of an attack. The expected utility for arbitrary $p(\psi|\chi)$ and P_{w} is obtained from (1). The objective of the policymaker, for a given R_c, is to choose the defense expenditures $\hat{\chi}(R_c)$ for which attains an expected reward that does not fall below R_c for the widest horizon of uncertainty.

In order to find the robust-optimal allocation we first need to express the robustness to model uncertainty ($\widehat{\alpha}$) as a function of the defense allocation. Full details of this representation are provided in the Appendix. Intuitively, we do that by exploiting the assumption that marginal utility is positive and decreasing in ex-post resources, and choose probability objects, to be denoted by $\hat{P}_w$ and $\hat{p}$, which are 'worse' than $\widetilde{p}(\psi|\chi)$ and $\widetilde{P}_{\mathrm{w}}(\chi)$, respectively, but differ from these proportionally by no more than α%. Since an attack can only decrease expected welfare, the probability of attack is taken to be:

$$\hat{P}_w = \begin{cases} (1+\alpha)\widetilde{P}_{\mathrm{w}} & \text{if } \alpha < (1-\widetilde{P}_{\mathrm{w}})/\widetilde{P}_{\mathrm{w}} \\ 1 & \text{else} \end{cases} \tag{12}$$

For the damage pdf, we distinguish between $\alpha \leq 1$, and $\alpha > 1$. For the former we use:

$$\hat{p}(\psi|\chi) = \begin{cases} (1-\alpha)\widetilde{p}(\psi|\chi) & \text{if } \psi \leq \psi_{\mathrm{m}} \\ (1+\alpha)\widetilde{p}(\psi|\chi) & \text{else} \end{cases} \tag{13}$$

where ψ_{m} is the median of the best estimate of the damage pdf $\widetilde{p}(\psi|\chi)$. For $\alpha > 1$ we use:

$$\hat{p}(\psi|\chi) = \begin{cases} 0 & \text{if } \psi \leq \psi_{\mathrm{s}} \\ (1+\alpha)\widetilde{p}(\psi|\chi) & \text{else} \end{cases} \tag{14}$$

where ψ_{s} satisfies:

$$(1+\alpha)\int_{\psi_{\mathrm{s}}}^1 \widetilde{p}(\psi|\chi)\,\mathrm{d}\psi = 1 \tag{15}$$

In other words, ψ_{s} is the $1-1/(1+\alpha)$ quantile of $\widetilde{p}(\psi|\chi)$.

Next we define:

$$\widetilde{r}_1(\chi) \quad = \quad \int_0^{\psi_{\mathrm{m}}} u\left[(1-\psi)(1-\chi_1-\chi_2)\right]\widetilde{p}(\psi|\chi)\,\mathrm{d}\psi \tag{16}$$

$$\widetilde{r}_2(\chi) \;=\; \int_{\psi_{\rm m}}^{1} u\left[(1-\psi)(1-\chi_1-\chi_2)\right]\widetilde{p}(\psi|\chi)\,\mathrm{d}\psi \tag{17}$$

$$\widetilde{r}(\chi) \;=\; \widetilde{r}_1+\widetilde{r}_2 \tag{18}$$

$$\delta_r(\chi) \;=\; \widetilde{r}_1-\widetilde{r}_2 \tag{19}$$

$$c(\chi) \;=\; 1-\chi_1-\chi_2 \tag{20}$$

Since marginal utility is positive and decreasing, $\widetilde{r}_1 > \widetilde{r}_2$ and $\widetilde{p}(\psi \leq \psi_{\rm m})$=$\widetilde{p}(\psi > \psi_{\rm m})$=0.5. Consequently, as we show in the Appendix, the smallest expected utility over $\mathcal{F}(\alpha, \widetilde{p}, \widetilde{P}_{\rm w})$, is obtained with $\hat{P}_w$ and $\hat{p}$, is given by (33) and (34) for $\alpha \leq 1$ and $\alpha > 1$, respectively. Equating these terms to R_c, and solving these quadratic expressions for α yields the desired robustness as a function of the defense allocation chosen. Suppressing the explicit dependence on χ, the functional form of $\widehat{\alpha}$ that corresponds to (13) is:

$$\widehat{\alpha}(\chi|R_c) = \begin{cases} \dfrac{(2\widetilde{r}_2-u(c))\widetilde{P}_{\rm w}-\sqrt{\left(2\widetilde{r}_2-u(c)\right)^2\widetilde{P}_{\rm w}^2+4\delta_r\widetilde{P}_{\rm w}\left((\tilde{r}-u(c))\widetilde{P}_{\rm w}+u(c)-R_c\right)}}{2\delta_r\widetilde{P}_{\rm w}} & \text{if } \widetilde{R}(\chi)\geq R_{\rm c} \\ 0 & \text{else} \end{cases}, \tag{21}$$

and the functional form of $\widehat{\alpha}$ corresponding to (14) is:

$$\widehat{\alpha}(\chi|R_{\rm c}) = \begin{cases} \dfrac{(2\widetilde{s}_2-u(c))\widetilde{P}_{\rm w}+\sqrt{\left(2\widetilde{s}_2-u(c)\right)^2\widetilde{P}_{\rm w}^2-4\widetilde{s}_2\widetilde{P}_{\rm w}\left((\widetilde{s}_2-u(c))\widetilde{P}_{\rm w}+u(c)-R_c\right)}}{-2\widetilde{s}_2\widetilde{P}_{\rm w}} & \text{if } \widetilde{R}(\chi)\geq R_{\rm c} \\ 0 & \text{else} \end{cases}, \tag{22}$$

where:

$$\widetilde{s}_2 \;=\; \int_{\psi_{\rm s}}^{1} u\left[(1-\psi)(1-\chi_1-\chi_2)\right]\widetilde{p}(\psi|\chi)\,\mathrm{d}\psi \tag{23}$$

The value of $\widetilde{R}$ in (11) is expressed in terms of utility. In our economy the representative individual has a maximum of a single unit of consumption, from which defense expenditures and any damage to the economy are deducted. Define the allocation that maximizes expected utility under the best estimates of an attack probability and damages pdf:

$$\chi^* \equiv \arg\max_{\chi} \tilde{R}\left(\chi|\widetilde{p}(\cdot|\chi), \widetilde{P}_{\rm w}(\chi)\right). \tag{24}$$

We use this allocation to generate a loss function in terms of a compensating consumption differential q. For that purpose define $L(q)$ to be the expected utility under $\widetilde{p}$ and $\tilde{P}_w$ obtained from defense allocation χ^* when an amount q of consumption is lost in addition to $\sum_i \chi_i^*$:

$$L(q) = \widetilde{P}_{\rm w}(\chi^*)\int_0^1 u\left((1-\psi)(1-\chi_1^*-\chi_2^*-q)\right)\widetilde{p}(\psi|\chi^*)\,d\psi + \left(1-\widetilde{P}_{\rm w}(\chi^*)\right)u(1-\chi_1^*-\chi_2^*-q) \tag{25}$$

In the numerical results below, rather than choosing values of R_c or reporting the values of $\widetilde{R}(\chi)$ we choose to express losses in terms of the consumption equivalents. Hence the maximum acceptable losses are expressed in terms of the consumption equivalents q_c, so that deviations from maximized expected utility is $q_c = L^{-1}(R_c)$. In similar fashion, the expected utility $\widetilde{R}(\chi)$ associated with a particular defense allocation χ, expressed in consumption equivalences is $\widetilde{q} = L^{-1}\left(\widetilde{R}(\chi)\right)$, where $L^{-1}\left(\widetilde{R}(\chi^*)\right) = 0$ by definition.

3.4 Historical Perspectives

In terms of the damage inflicted on society by war, Table 1 presents the direct expenses of major U.S. wars along with U.S. fatalities. In proportion to the size of the economy, the heaviest direct economic costs were incurred during World War II; however, in terms of loss of life relative to the size of the population, the Civil War was far deadlier. Obviously the latter conflict, fought entirely on American soil, generated the greatest loss of capital, damage to infrastructure and disruption of output. In the aftermath of World War I, Bogart (1919) began developing the tools to measure, compare, and aggregate all these different costs of war, by including not only the direct costs in military expenditure as well as physical destruction, but the indirect costs associated with the capitalized values of losses in life and lost production. According to his calculations, the share of direct costs incurred by all the combatants of World War I amounted to only 55% of the total losses the war generated.

Broadberry and Howlett (1998) calculate that the UK spent approximately half of its GDP fighting World War II during the years 1940 to 1944. In addition it suffered losses of physical capital that amounted to 89% of GDP in 1938 (see Mitchell, 1980), and human capital losses (calculated conservatively in terms of just the schooling invested in those killed) of 2.5% of GDP in 1938. By any measure Soviet losses were far higher. During 1942 and 1943 defense expenditures in the Soviet Union reached 61% of GDP, losses of physical capital amounted to 223% of pre-war GDP and losses of human capital were 109% (Harrison, 1998). Of course these figures do not include the extraordinary privations suffered by those living under German occupation during much of this period. To study the cost of World War II for the United States, Rockoff (1998) employs a counterfactual approach developed by Goldin and Lewis (1975) to study the US Civil War. According to his estimates the total present value of foregone consumption that can be attributed to both direct and indirect losses generated by the war equals 2.27 years of consumption in 1941.[8]

[8] As most of its effects are indirect, the impact of terror is not as well understood. Abadie and Gardeazabal (2003) estimate that terrorism in the Basque country of Northern Spain has reduced GDP by 10%. Similarly, recent estimates of the loss in GDP that can be attributed to three years of recurring terrorism against Israel is also approximately 10% (Eckstein and Tsiddon, 2004 and Persitz, 2005). Bram, Orr and Rapaport (2002) estimate the total cost of one incident, the September 11, 2001, attack on the World Trade Center in New York, including lost lifetime earnings of those killed, to be between $33 billion and $36 billion. What is clear is that large-scale conventional warfare is far more costly than any losses associated with terror—Hess (2003) calculates

Conflict	Total Direct Costs		People Mobilized		Fatalities	
	Billions of 2002$	*Percent of GDP*	*Thousands*	*Percent of Pop.*	*Numbers*	*Percent of Pop.*
Revolutionary War (1775 – 1783)	2.2	63%	200	5.70%	4,435	0.127%
War of 1812 (1812 – 1815)	1.1	13%	286	3.80%	2,260	0.030%
Mexican War (1846 – 1848)	1.6	3%	79	0.40%	1,733	0.008%
Civil War (1861 – 1865)	62	104%	3,868	11.10%	184,594	0.538%
Span. Amer. War (1898)	9.6	3%	307	0.40%	385	0.001%
World War I (1917 – 1918)	190.6	24%	4,744	4.60%	53,513	0.052%
World War II (1941 – 1945)	2,896.3	130%	16,354	12.20%	292,131	0.219%
Korea (1950 – 1953)	335.9	15%	5,764	3.80%	33,651	0.022%
Vietnam (1964 – 1972)	494.3	12%	8,744	4.30%	47,369	0.023%
First Gulf War (1990 – 1991)	76.1	1%	2,750	1.10%	148	0.000%

Table 1: American Costs and Casualties from Major Wars. *Source*: William D. Nordhaus, "The Economic Consequences of a War with Iraq" in *War with Iraq*, ed. Kaysen, *et. al.*

3.5 Example Specifications of the Policymaker's Best Estimates of Model Components

This subsection presents the functional forms that we use in our subsequent illustration. These specifications were chosen to reflect the historical magnitudes described above, and to capture the intertwined channels through which defense allocations affect both the probability of suffering an attack and the losses such an attack would impose. The best estimate of the pdf of damage, conditional on being attacked, is assumed to be the Beta function:

$$\widetilde{p}(\psi|\chi) = \frac{\psi^{a(\chi)-1}(1-\psi)^{b(\chi)-1}\Gamma(a(\chi)+b(\chi))}{\Gamma(a(\chi))\Gamma(b(\chi))}, \tag{26}$$

that for countries that have experienced conflict between 1960-1992 (nearly all of it civil war or terrorism in this period), the loss in welfare associated with these conflicts is on average equivalent to a permanent 8% drop in their consumption.

where the Gamma function is given by $\Gamma(x) = \int_0^\infty t^{(x-1)}e^{-t}dt$, z is the maximum amount of damage an attack can create as a fraction of GDP, and:

$$a(\chi) = 1 + \frac{\chi_1}{\chi_2} + \theta e^{\theta\chi_1} \ln(\chi_1 + \chi_2)\,\chi_1, \tag{27}$$

$$b(\chi) = 2 + \frac{\chi_1}{\chi_2} + e^{\chi_2} \ln(\chi_1 + \chi_2)\,\chi_2. \tag{28}$$

The mean of this pdf is:

$$\mathrm{E}(\psi|\chi) = \frac{\Gamma(1 + a(\chi))\Gamma(a(\chi) + b(\chi))}{\Gamma(a(\chi))\Gamma(1 + a(\chi) + b(\chi))} \tag{29}$$

We set the value of $\theta = 4$ in our calculations.

The best estimate of the probability of suffering an attack is:

$$\widetilde{P}_{\mathrm{w}}(\chi) = 1 - \left(\frac{\beta_{\frac{1}{2}}(a(\chi), b(\chi))\Gamma(a(\chi) + b(\chi))(\chi_1 + \chi_2)}{\Gamma(a(\chi))\Gamma(b(\chi))}\right)^{\frac{1}{5}} \tag{30}$$

where $\beta_x(a, b) \equiv \int_0^x t^{a-1}(1 - t)^{b-1}dt$ and $a(\chi)$ and $b(\chi)$ are defined in (27) and (28). The term $\chi_1 + \chi_2$ represents the total military expenditure, and reflects its deterrent value. In our example there is an inverse relationship between the expected damage (conditioned on being attacked) that corresponds to a particular allocation of defense spending χ and the probability associated with that allocation, of being attacked in the first place. This can be seen in the term $\frac{\beta_{\frac{1}{2}}(a(\chi),b(\chi))\Gamma(a(\chi)+b(\chi))(\chi_1+\chi_2)}{\Gamma(a(\chi))\Gamma(b(\chi))}$ in (30) and represents the probability that damages will be small—less than half the maximal possible damage—conditioned on attack taking place. Hence, the higher this number is, the lower the likelihood that an adversary will be tempted to launch an attack. Finally, we adopt the constant-risk-aversion specification for the utility function $u(c) = c^{1-\gamma}/(1 - \gamma)$, where $\gamma > 0$.[9]

3.6 Numerical Results

We focus on a scenario with relatively high values of potential damage in the event of attack, and choose parameters accordingly. For example, setting the overall size of the defense budget to equal 10% of total available income, equally divided between χ_1 and χ_2, implies given the functional specifications in (26) and (30), that expected losses in the event of attack will amount to an equivalent loss of 33% of annual consumption, where the probability of an attack taking place is just under 40%. Doubling the resources devoted to each, leaves the conditional loss equal to 27% and reduces the probability of attack to just under 30%, and doubling the resources yet again still leaves an expected loss of nearly 12% conditional on being attacked with an attack probability of just under 23%. Overall, given the functional specifications in (26) and (30)

[9] In our calculations we set $\gamma = .98$, so utility is close to logarithmic.

of the probability of an attack and the damage density as functions of the defense allocation, policy makers who maximize expected welfare under (26) and (30) will devote just under 0.0842 of available resources to defense, with slightly more than half that, (χ_1=0.0465), allocated to RMA-type expenditure and the remainder, (χ_2=0.0376), to traditional military spending. Yet what is important to emphasise is that these probabilities of attack and damage density functions are merely best estimates, and the allocation that maximises expected utility is not robust to any degree of model uncertainty.

In Figure 3 we demonstrate, for alternative levels of the lowest acceptable expected utility R_c, expressed in consumption losses equivalent units q_c=0.05, 0.1, 0.15 and 0.2, the allocations χ_1 and χ_2 that attain maximum robustness to model uncertainty at the values $\widehat{\alpha}$=0.179, 0.392, 0.660 and 1.071, respectively. The contours in each frame represent the trade-offs between χ_1 and χ_2 for the fixed value of q_c of that frame and a degree of robustness $\widehat{\alpha}$ marked next to the contour.

Consider the first panel in Figure 3 which corresponds to the consumption equivalent loss of q_c=.05. Calculating the value of $\widehat{\alpha}$ for different combinations of χ_1 and χ_2 in (21) and (22) yields the contours, each one corresponding to a different value of $\widehat{\alpha}$ that surround the unique maximum at $\widehat{\alpha}$=0.179. That maximum is attained at χ_1=0.0583 and χ_2=0.0661. This means that a policy maker willing to accept utility equivalent of 5% less consumption compared to that associated with expected utility maximisation can, by slightly increasing expenditure on RMA-type weaponry from χ_1=0.0465 to χ_1=0.0583, but nearly doubling expenditure on traditional military hardware from χ_2=0.0376 to χ_2=0.0661, ensure at least 95% of the (consumption equivalent) expected utility maximizing reward even if the best estimates of both the likelihood of attack, and the conditional damage are off by as much as 17.9%.

The remaining three panels in Figure 3 further demonstrate the trade-offs between how much a policy maker is prepared to accept a lower level of utility (and hence higher values of q_c), and the maximum degrees of robustness to model uncertainty that this sacrifice affords. Thus by setting q_c=0.1, the degree of maximum robustness rises to 39.2%, achieved by setting χ_1=0.0717 and χ_2=0.1026; setting q_c=0.15, the degree of maximum robustness rises to 66.0%, achieved by setting χ_1=0.0882 and χ_2=0.1496; and setting q_c=0.2, the degree of maximum robustness rises to 107.1%, achieved by setting χ_1=0.114 and χ_2=0.2209. In each case higher robustness is achieved by allocating more resources to defense spending. This is hardly surprising. What is being chosen here are allocations of defense expenditure designed to insure against model misspecification that is too optimistic. Yet the increases are *asymmetric*—relatively little of the additional expenditure is allocated to the new RMA-type weaponry. Instead higher levels of robustness are mostly achieved by sharply increasing the investment in the more traditional weaponry.

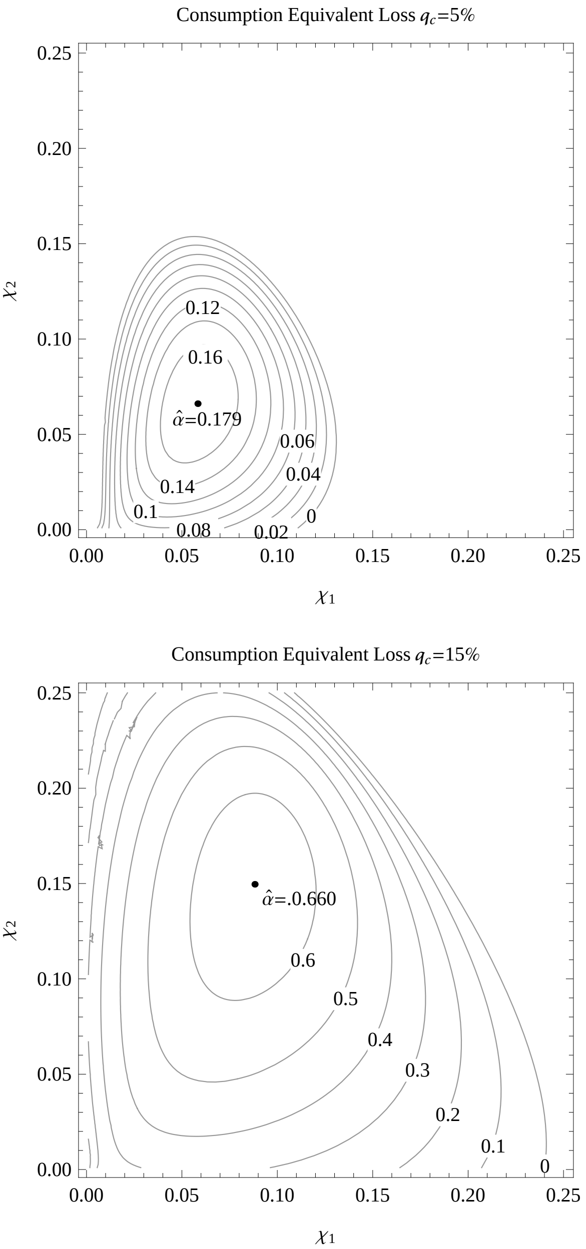

Figure 3: Contour plots for different combinations of χ_1 and χ_2. The contours represent different values of $\widehat{\alpha}$ where minimum value of acceptable expected utility is equivalent to a loss of 5%, 10%, 15% or 20% relative to maximised expected utility.

Why should the share of resources devoted to traditional war fighting capacity rise so much faster than RMA-type expenditures as the desired degree of robustness to model uncertainty increases? The reason is that while increasing both types of security expenditure may generate higher levels of deterrence and also lower the amount of damage suffered in the event of an attack, it is the acquisition of planes, tanks and artillery, rather than high precision weaponry and enhanced intelligence gathering that works best in lowering the likelihood of extreme damage. This pattern can be seen in Figure 4 which traces the the combinations of χ_1 and χ_2 that yield maximum values of robustness and the contours reflect the losses in actual expected utility in consumption terms $\widetilde{q}$.

Table 2 summarises the results in Figure 3. It emphasises the relationship that exists between maximum robustness, $\widehat{\alpha}$ and the acceptable level of losses q_c, and how they are achieved by different combinations of the two types of expenditure. The last column in Table 2 lists the values of $\widetilde{q}$ that express how much in consumption terms we actually expect to relinquish (rather than q_c which represents how much we are prepared to relinquish) relative to expected utility maximisation under the best estimates $\widetilde{P}_{\mathrm{w}}$ and $\widetilde{p}$. That is, while the maximum loss of expected utility will be less than q_c for all eventualities in a $\widehat{\alpha}$ surrounding of $\widetilde{P}_{\mathrm{w}}$ and $\widetilde{p}$ corresponding to any particular row in Table 2, the expected loss of utility relative to maximised expected utility is only $\widetilde{q}$.

What is important to emphasise is that whereas the manner in which an expected utility maximiser allocates expenditure pays little attention to the perceived low probability events associated with extreme losses, a robust satisficer, wary of relying too heavily on the probabilities themselves, will behave in a way that both favours higher expenditure, with the lion's share of additional resource devoted to the type of spending that is less susceptible to catastrophic failure. This illustrates how the search for robustness alters decision making when the trade-offs satisfy (7) and (8).

To better understand the implicit tradeoff between robustness and minimum acceptable levels of utility, we plot in Figure 5 the implicit relationship between $\widehat{\alpha}$ and $\widetilde{q}$ for combinations of χ_1 and χ_2 between 0 and 0.25 (in increments of 0.01). Each curve represents a particular combination of defense expenditure (χ_1, χ_2) and traces the implicit relationships in (21) and (22) between different values of $\widehat{\alpha}$ and $\widetilde{q}$ (with the aid of the loss function (25)). Each curve corresponds to a particular defense allocation, and begins on the horizontal axis where $\widehat{\alpha}$=0 and the welfare loss $\widetilde{q}$ is the one implied by the deviation of that allocation from χ^*. As we consider larger deviations from the best guess uncertainty model, (i.e. increase $\widehat{\alpha}$), that particular allocation yields larger welfare losses, and hence the curves are upward sloping. The upper contours of all the grey lines generate the envelope that defines the relationship between maximum robustness $\widehat{\alpha}$ and $\widetilde{q}$, starting at the origin which corresponds to the policy that generates maximum expected utility

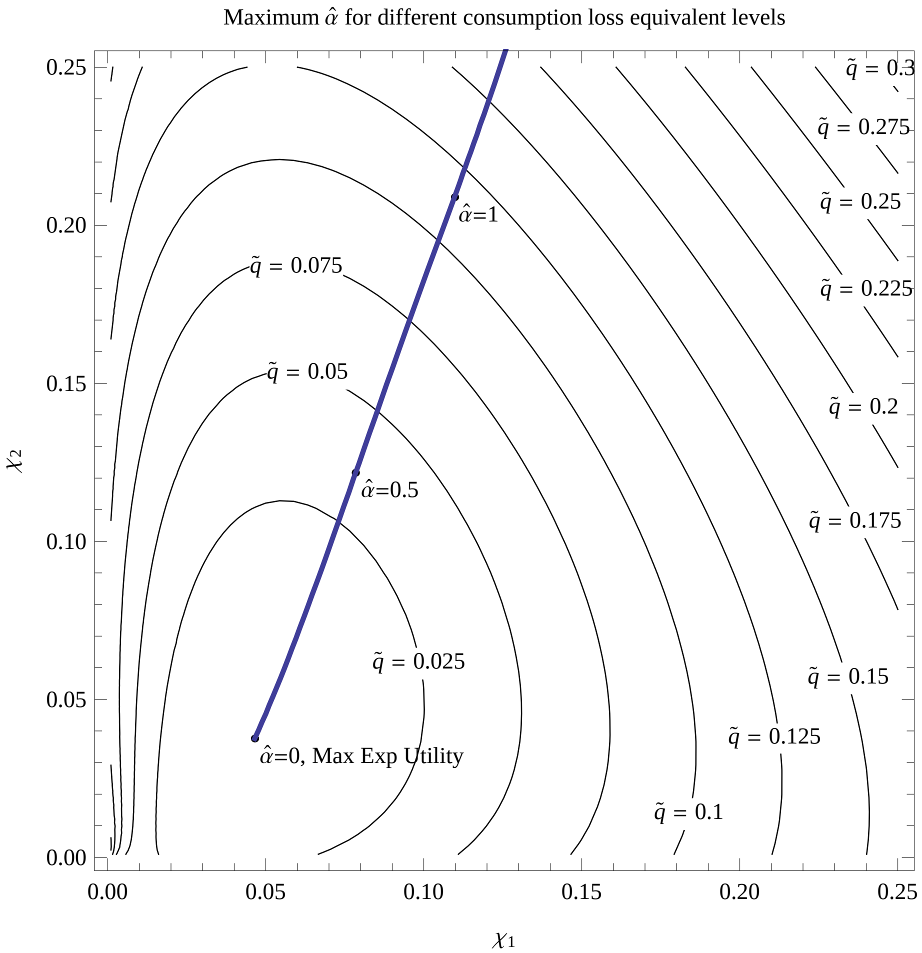

Figure 4: Contour plots for different combinations of χ_1 and χ_2. The contours represent different values of $\widetilde{q}$, the consumption equivalent expected loss in utility relative to maximised expected utility. The solid thick curve represents all the combinations of χ_1 and χ_2 attaining maximum robustness $\widehat{\alpha}$ to model misspecification.

and zero robustness.

Policy makers first determine the minimum level of acceptable utility, denominated here as consumption equivalent deviations from maximum expected utility, and this determines the maximum degree of attainable robustness to model error. A given combination of defense expenditure (χ_1, χ_2) delivers a monotonic relationship between the value of q_c and $\widehat{\alpha}$ in Figure 4, which together generate an envelope that defines the convex relationship between q_c and $\widehat{\alpha}$ (the red curves correspond to the particular values of χ_1 and χ_2 that generate maximum robustness for the values $q_c = 0.05, 0.1, 0.15, 0.2$). The cost of achieving an additional increment of robustness to model uncertainty and immunity to catastrophic failure declines as policy makers permit greater deviations from policies that maximize expected utility. The more policy makers are cognizant of the limitations of the information they possess, in terms of efficacy of the measures they have at their disposal to both deter and repel aggression, the more we

q_c	χ_1	χ_2	$\hat{\alpha}$	$\widetilde{q}$
0	0.0465	0.0376	0	0
0.05	0.0583	0.0661	0.1794	0.0053
0.1	0.0717	0.1026	0.3919	0.0228
0.15	0.0882	0.1496	0.6601	0.0569
0.2	0.1140	0.2209	1.0712	0.1283

Table 2: The combinations of χ_1 and χ_2 that generate maximum robustness $\widehat{\alpha}$, for different levels of maximally acceptable consumption equivalent losses compared to expected utility maximization, q_c, and the implied loss in expected utility (in consumption equivalents) of these defense allocations $\widetilde{q}$.

should expect them to reach for solutions that achieve robustness through both greater military expenditure and heavier reliance on the type of expenditure that is most effective at preventing catastrophic losses.

3.7 Discussion

How do we quantify security threats and the losses they may generate, and what are the different possible security expenditures that are meant to counter them? Fragmentary anecdotal and quantitative evidence exists concerning the impact of military expenditures on the probability distribution of war-related damage. Rohlfs (2006) estimates that the marginal effectiveness of a U.S. tank in Western Europe for 164 battles during World War II was twenty-four times the effectiveness of a single infantryman but eighty-seven times as expensive to use. The discrepancy is explained by the higher casualties associated with intensive use of infantry, implying that the U.S. government assigned a value of approximately one million dollars (in 2003 dollars) to each soldier's life saved on the battlefield. Although Rohlfs' study involves calculating ex-post a relatively simple trade-off in a single theater of a war in its fifth and sixth year (1944-1945), his estimates contain relatively large standard errors and vary across different sub-samples. By contrast, policy makers must determine ex-ante both the overall effectiveness of defense expenditures and their optimal allocations, at a stage when the nature, scale, and even eventuality of a conflict may be only hypothetical. In the years prior to being attacked in 1941, the U.S.S.R. engaged in a massive rearmament program that, according to Bergson (1961), lowered per-capita consumption between 1937 and 1940 by as much as 8.4%. However, not only did rearmament fail to deter Hitler's invasion as the Soviets had hoped, but because of serious military miscalculations, much of the arms and manpower was squandered during the summer and fall of 1941 without seriously slowing the German advance (Harrison, 1985).

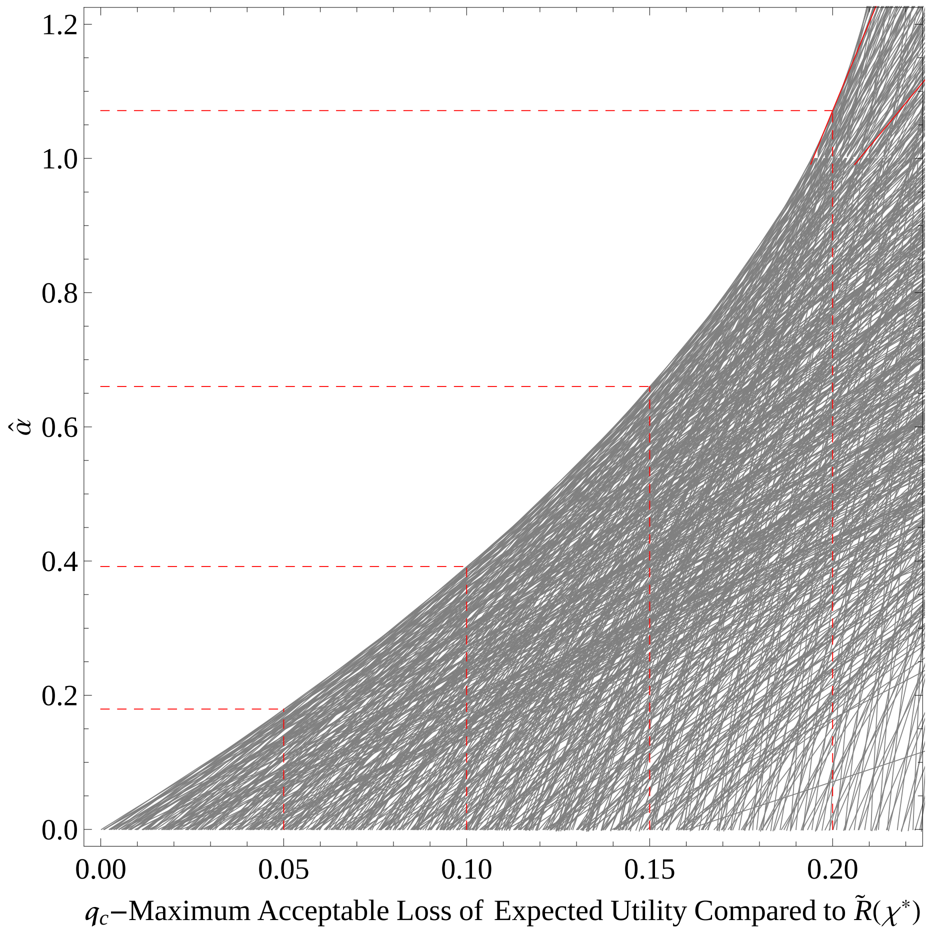

Figure 5: Each of the grey curves corresponds to a different combination of (χ_1, χ_2) and captures the tradeoff between $\widetilde{q}$ and $\widehat{\alpha}$. The envelope traced by the different grey curves represents the maximum acceptable loss expressed as the consumption equivalent q_c and the maximum attainable robustness $\hat{\alpha}$ to model misspecification.

Often enough, the pursuit of policies designed to limit casualties but not be robust to misspecification can result in complete catastrophe. Soon after the end of World War I the French government and military immediately began a debate on the best way to prepare for renewed conflict with Germany. One doctrine relied heavily on a smaller professional army well-trained in offense, and capable of mastering the new techniques of maneuver. The alternative that was eventually adopted, relied on a larger force of conscripts backed by reservists to defend France's borders along a static heavily fortified front. On 30 September 1927, the *Commission d'Organization des Région Fortifiées*, (CORF) was established to plan construction of a line of fortifications along France's western frontiers with Italy, Germany, portions of southern and central Belgium as well as Corsica, that ultimately came to bear the name of the Minister of War and decorated veteran of Verdun, André Maginot. The Maginot line was meant to protect France from a possible German invasion once France had completed the planned evacuation of its forces from the Rhineland in 1935. The line was based on a series of concrete multi-storied casements and *ouvrages* (fortresses). The smallest housed a squad of thirty men manning twin 7.5mm machine guns and 37 or 47mm anti-tank guns. The largest the *gros ouvrages* were large complex structures, built mostly underground, housing garrisons of 500-1,000 men and fired heavy artillery from turrets set in the ground. Each was surrounded by their own series of machine gun turrets for protection.

The Maginot line's appeal was its promise to both deter a German attack and in the event deterrence failed, to minimize French casualties by replacing 'a wall of chests' with 'a wall of concrete'. Between 1930 and the time most of it had been completed in 1935, its construction alone consumed between one-half and one percent of French output each year, at a time when total French expenditure on defense averaged 4.7% of GDP. [10] Two important gaps remained along the border with Belgium, near the sea the high water table precluded the building of subterranean forts, and in the Ardennes where Marshal Pétain, inspector general of the army until 1931, had asserted the terrain was impassible to German armour.[11]

During the six week German assault on France between 10 May 1940 and the signing of the armistice on 22 June, nearly every fort along the Maginot line performed as it was designed—they deterred direct attacks along the Franco-German border, and most successfully repulsed those efforts made by the Germans to attack them from the east. Ultimately however, Pétain's assurances not withstanding, the Germans did move their armour through the Ardennes, outflanking the Maginot line and forcing French capitulation. French strategic doctrine, designed to operate under one set of conditions, was not robust to the unexpected, and ultimately failed in a catastrophic manner. A doctrine that might have proven more robust to catastrophic failure—that advocated by General Charles De Gaulle (but successfully implemented by *Wehrmacht* General Heinz Guderian) of armoured forces able to move quickly to attack and fight independently of the infantry—was rejected, in part because it implied the likelihood of higher casualties.[12]

In our interpretation, investment in the Maginot line was then akin to choosing high values of χ_1 (RMA in our more modern example). Had the Germans not found a way to circumvent it, it might have repelled the invasion, and then done so as designed—at the cost of relatively few casualties. Instead it failed catastrophically—French planning was not sufficiently robust.

[10] W. J. Thies, "Alliances and Collective Goods," *Journal of Conflict Resolution*, 1987; 31: 298-332.

[11] J.E. Kaufmann and H.W. Kaufmann, *The Maginot Line: None Shall Pass*, Praeger Security International, Westport CT, 1997.

[12] J.E. Kaufmann and H.W. Kaufmann, *Fortress France: The Maginot Line and French Defenses in World War II*, Praeger Security International, Westport CT, 2006.

4 Conclusion

Our application of the information gap approach to the problem of allocating resources for national security yields three broad conclusions, each of which is evaluated quantitatively. First, the higher the robustness to model misspecification policy makers demand, the higher the overall level of defense expenditures required. Second, the greater the demand for robustness, the greater should be the reliance on those defense measures better suited to prevent extremely high levels of damage. Similarly, the more robustness policy makers wish to achieve, the more they should eschew investment in systems that promise the best *expected* outcomes on the battlefield, but are also more vulnerable to catastrophical damage when they fail.

Beyond the normative recommendations for how policy makers can best allocate resources to protect their countries from aggression, in a world in which much of the relevant information is unreliable, the model also provides some positive insights into the reasons for the policies policy makers choose today or have chosen in the past. In a world with unreliable probabilistic information, we should not be surprised if policy makers lavish higher expenditure on defense than would be appropriate if the only goal were expected utility maximization. Furthermore, we would expect policy makers to favor expenditures on weapons systems, and associated tactics and strategies, that are both most effective in preventing worst-case scenarios and are also better understood. These would suggest one possible rationale for military planners' reputation for conservatism, and indeed inertia, when confronted with new and untried technologies and doctrines.

A century before Frank Knight made the distinction between risk and uncertainty, Von Clausewitz noted the problems associated with the unreliability of information in military settings.

> Many intelligence reports in war are contradictory; even more are false, and most are uncertain .... One report tallies with another, confirms it, magnifies it, lends it color, till [the officer] has to make a quick decision which is soon recognized to be mistaken, just as the reports turn out to be lies, exaggerations, errors, and so on. In short, most intelligence is false, and the effect of fear is to multiply lies and inaccuracies ....[13] The general unreliability of all information presents a special problem in war: all action takes place, so to speak, in a kind of twilight, which like fog or moonlight, often tends to make things seem grotesque and larger than they really are.[14]

The information gap approach does nothing to ameliorate the unreliability of information, but

[13] Carl von Clausewitz, *On War*, trans. Michael Howard and Peter Paret, (Everyman's Library, London, 1993), 136.

[14] *Ibid.*,161.

it does offer policy makers a methodology to make decisions better suited for such environments.

5 Appendix-Derivation of the Robustness Function

We derive the robustness function in (21) for values of the robustness not in excess of unity: $\widehat{\alpha} \le 1$. We make no assumptions about the utility function $u(c)$ other than that the marginal utility is positive: $u'(c) > 0$.

The main task is to find the pdf of the damage, $p(\psi|\chi)$ which, at horizon of uncertainty α, minimizes the expected utility $R(\chi|p, P_{\rm w})$ defined in (1). Because the marginal utility is positive it is evident that $R(\chi|p, P_{\rm w})$ is minimized by that pdf in $\mathcal{F}(\alpha, \widetilde{p}, \widetilde{P}_{\rm w})$ which assigns as much weight as possible at large levels of damage and as little weight as possible at low levels of damage. For the fractional-error info-gap model in (9) one readily shows that $\min\limits_{\alpha} R(\chi|p, P_{\rm w})$ occurs with the following pdf:

$$p(\psi|\chi) = \begin{cases} (1-\alpha)\widetilde{p}(\psi|\chi) & \text{if } \psi \le \psi_{\rm m} \\ (1+\alpha)\widetilde{p}(\psi|\chi) & \text{else} \end{cases} \tag{31}$$

where $\psi_{\rm m}$ is the median of the estimated pdf $\widetilde{p}(\psi|\chi)$ and where $\alpha \le 1$.

If $\alpha > 1$ then $\min\limits_{\alpha} R(\chi|p, P_{\rm w})$ occurs with the following pdf:

$$p(\psi|\chi) = \begin{cases} 0 & \text{if } \psi \le \psi_{\rm s} \\ (1+\alpha)\widetilde{p}(\psi|\chi) & \text{else} \end{cases} \tag{32}$$

where $\psi_{\rm s}$ satisfies:

$$(1+\alpha)\int_{\psi_{\rm s}}^{1} \widetilde{p}(\psi|\chi)\, {\rm d}\psi = 1 \tag{33}$$

In other words, $\psi_{\rm s}$ is the $1 - 1/(1+\alpha)$ quantile of $\widetilde{p}(\psi|\chi)$.

Consider first the case α<1.

The utility $R(\chi|p, P_{\rm w})$ in (1), evaluated with the pdf in (31), is:

$$R(\chi|p, P_{\rm w}) = (\widetilde{r} - \delta_r\alpha - u(c))\, P_{\rm w} + u(c) \tag{34}$$

where $\widetilde{r}$ and δ_r are defined in (10) and (19) and $u(c)$ is the utility if an attack does not occur. The term $\widetilde{r} - \delta_r\alpha - u(c)$ is negative so the minimizing value of $P_{\rm w}$ in $\mathcal{F}(\alpha, \widetilde{p}, \widetilde{P}_{\rm w})$ is $(1+\alpha)\widetilde{P}_{\rm w}$. Thus the minimum expected utility, up to horizon of uncertainty α, is:

$$\min_{p, P_{\rm w} \in \mathcal{F}(\alpha, \widetilde{p}, \widetilde{P}_{\rm w})} R(\chi|p, P_{\rm w}) = (\widetilde{r} - \delta_r\alpha - u(c))\,(1+\alpha)\widetilde{P}_{\rm w} + u(c) \tag{35}$$

Given that marginal utility is positive, $\widetilde{r}_1 > \widetilde{r}_2$, and therefore $\delta_r > 0$. This, together with the assumption $\widetilde{r}(\chi) < u(c)$, implies that $\widetilde{r} - u(c) - \delta_r < 0$. Denote the minimum value in (35) as $R_{\rm c}$ and solve the quadratic function in α for the positive root (21).

For instances where $\alpha \geq 1$:

$$\min_{p, P_{\mathrm{w}} \in \mathcal{F}(\alpha, \widetilde{p}, \widetilde{P}_{\mathrm{w}})} R(\chi | p, P_{\mathrm{w}}) = ((1+\alpha)\widetilde{s}_2 - u(c))\,(1+\alpha)\widetilde{P}_{\mathrm{w}} + u(c) \tag{36}$$

and following the same procedure yields the solution for the positive root α in (22).

References

[1] A. Abadie and J. Gardeazabal, "The Economic Costs of Conflict: A Case Study of the Basque Country," *American Economic Association*, 2003; 93: 113-132.

[2] Y. Ben-Haim, *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*, Academic Press, London, 2nd edition, 2006.

[3] Y. Ben-Haim, Info-gap Decision Theory For Engineering Design. Or: Why 'Good' is Preferable to 'Best', in E. Nikolaidis, D. Ghiocel and S. Singhal, eds., *Engineering Design Reliability Handbook*, CRC Press, Boca Raton, 2005.

[4] A. Bergson, *The Real National Income of Soviet Russia Since 1928*, Harvard University Press, Cambridge, MA, 1961.

[5] E. L. Bogart, *The Direct and Indirect Costs of the Great War*, Oxford University Press, New York, 1919.

[6] J. Bram, J. Orr, and C. Rapaport, "Measuring the Effects of the September 11 Attack on New York City," *Federal Reserve Bank of New York Economic Policy Review*, 2002; 8: 5-20.

[7] S. Broadberry and P. Howlett, The United Kingdom: Victory at all Costs, in M. Harrison, ed., *The Economics of World War II*, Cambridge University Press, Cambridge, UK, 1998.

[8] R. J. C. Butow, *Tojo and the Coming of the War*, Stanford University Press, CA, 1961.

[9] E. A. Cohen, "A Revolution in Warfare," *Foreign Affairs*, 1996; 75: 37-54.

[10] A. J. Echevaria II, Enduring Elements of Military Strategy, in P. M. Cronin ed., *The Impenetrable Fog of War*, Praeger Security International, London, 2008.

[11] Z. Eckstein and D. Tsiddon, "Macroeconomic Consequences of Terror: Theory and the Case of Israel," *Journal of Monetary Economics*, 2004; 51: 971-1002.

[12] A. I. George, *Presidential Decisionmaking in Foreign Policy: The Effective Use of Information and Advice*, Westview Press, Boulder, CO., 1980.

[13] C. D. Goldin and F. D. Lewis, "The Economic Cost of the American Civil War: Estimates and Implications," *Journal of Economic History*, 1975; 35: 299-326.

[14] M. I. Handel, *War, Strategy and Intelligence*, Frank Cass, London, 1989.

[15] M. Harrison, *Soviet Planning in Peace and War*, Cambridge University Press, Cambridge, UK, 1985.

[16] G. D. Hess, "The Economic Welfare Cost of Conflict: An Empirical Assessment," Working Paper, Clarement Colleges, 2003.

[17] D. Hume, *An Enquiry Concerning Human Understanding*, Dover Publications Inc., Mineola, NY, 2004.

[18] T. L. Hughes, *The Fate of Facts in a World of Men*, Headline Series, No. 233, December 1976, Foreign Policy Association, Inc., New York.

[19] J.E. Kaufmann and H.W. Kaufmann, *The Maginot Line: None Shall Pass*, Praeger Security International, Westport CT, 1997.

[20] J.E. Kaufmann and H.W. Kaufmann, *Fortress France: The Maginot Line and French Defenses in World War II*, Praeger Security International, Westport CT, 2006.

[21] B. R. Mitchell, *European Historical Statistics, 1750-1975*, Macmillian Press Ltd., London, 1980.

[22] W. D. Nordhaus, The Economic Consequences of a War with Iraq, in C. Kaysen, S.E. Miller, M.B. Miller, W.D. Nordhaus and J.D. Steinbruner, eds., *War with Iraq: Costs, Consequences, and Alternatives*, American Academy of Arts and Sciences, Cambridge, MA, 2002.

[23] D. Persitz, "Counterfactual Analysis of Terrorism Effect on the Economy: the Case of Israel," Working Paper, Tel-Aviv University, 2005.

[24] H. Rockoff, The United States: From Ploughshares to Swords, in M. Harrison, ed., *The Economics of World War II*, Cambridge University Press, Cambridge, UK, 1998.

[25] C. A. Rohlfs, "The governments valuation of life-saving in war: a cost minimization approach," *Papers and Proceedings of the American Economic Association*, 2006; 96: 39-44.

[26] W. J. Thies, “Alliances and Collective Goods,” *Journal of Conflict Resolution*, 1987; 31: 298-332.

[27] C. von Clausewitz, *On War*, trans. Michael Howard and Peter Paret, Everyman’s Library, London, 1993.