

**Performance Measurement in Military Operations:
Information versus Incentives**

[Running title: "Measurement in Military Operations"]

Leo J. Blanken

Department of Defense Analysis

Naval Postgraduate School

Monterey, CA 93943

Jason J. Lepore

Department of Economics

Orfalea College of Business

California Polytechnic State University

San Luis Obispo, CA 93407

Abstract

We explore the impact of strategic assessment efforts on military organizations at war. To do so, we construct a model to explore the impact of a principal's choice among imperfect performance metrics for a military operation. In doing so, the principal must consider both the incentivizing and informational properties of the metric. We show the conditions under which uncertainty regarding the nature of the agent, as well as uncertainty regarding the operational environment, drives a metric choice that induces pathological behavior from the agent. More specifically, a poor metric choice can create an overly optimistic assessment and end up prolonging the conflict. We illustrate the model's insights in the cases of the Second World War and the Vietnam War.

INTRODUCTION

Assessing progress in war is always a difficult endeavor (Gartner 1997), but as Campbell, O'Hanlon, and Shapiro (2009: 16) argue: "counterinsurgency and stabilization operations --- like the ones in Iraq and Afghanistan --- are different, and more complex... How do we measure progress in such situations? This question is crucially important. Only by tracking progress can we know whether a strategy is working." Such relevant information may include the status of enemy forces, the performance of one's own forces, the local civilian population, and similar factors. The implication is that, with better assessments of the conflict, efforts of the military can be more efficiently utilized and victory will be more likely.

Many critics of current war efforts have complained that the measurement problem stems fundamentally from a lack of information regarding the operational environment. The conclusion of one analysis of the military intelligence efforts of the war in Afghanistan, for example, is damning on this topic:

Eight years into the war in Afghanistan... the vast intelligence apparatus is unable to answer fundamental questions about the environment... Ignorant of local economics and landowners, hazy about who the powerbrokers are and how they might be influenced, incurious about the correlations between various development projects and the levels of cooperation among villagers, and disengaged from the people in the best position to find answers... US intelligence officers and analysts can do little but shrug... (Flynn et al. 2010: 7).

In contrast to such pleas for more emphasis on the informational content of wartime measurement, however, we engage an under-explored ramification of the assessment problem, namely the impact of the selected metrics on the incentives of the military agent. In doing so, we examine an aspect of metrics that has been almost entirely neglected in the current policy debates: the manner in which one measures progress incentivizes the behavior of those who are conducting the war.

In his classic article, “On the Folly of Rewarding A, While hoping for B,” Steven Kerr (1975) introduced this problem in the wartime setting. He did so by recognizing that incentives matter even in motivating soldiers on the battlefield. In the Second World War the fact that personnel were drafted for the duration of the conflict motivated soldiers very differently than in case of Vietnam, in which soldiers were drafted for one-year tours of duty. Kerr explains how incentives were aligned in the Second World War.

What did the GI in World War II want? To go home. And when did he get to go home? When the war was won! If he disobeyed the orders to clean out trenches and take the hills, the war would not be won and he would not go home (1975: 771).

Kerr contrasts this with the incentives in the Vietnam conflict.

Consider the reward system in Vietnam. What did the soldier at the bottom want? To go home. And when did he get to go home? When his tour of duty was over! This was the case whether or not the war was won (1975: 771).

The important contribution of his work is two-fold: First, it points to the problem of incentive alignment as fundamental in military operations. Second, it makes clear that the structure of a particular conflict and the contracting of the soldiers can dictate war outcomes. We take Kerr’s insights and apply them to a formal treatment of the metric selection problem. We consider the possibility that the structure of the conflict and military contracting can be such that incentives between the principal and agent are aligned or misaligned; the key point we make is that the measurement problem is starkly different in the two settings.

To show this, we construct a principal-agent model with imperfect measurement of success and assumptions particular to the military context.¹ The principal is unable to observe (*ex ante* or *ex post*) actual operational effectiveness. Instead the principal must use an imperfect performance measure (or metric) to assess information about the current conflict. We consider two types of agents. The first type has incentives that are aligned

with the principal's and consequently its behavior does not respond to the metric. The second type of agent has misaligned incentives and looks to maximize the metric instead of actual operational success. In this case, the metric incentivizes the agent to choose its actions in a particular direction. We find that the informational content and incentive properties of a metric are dependent on different features of a metric. Information is based on the realizations of the metric, while the incentives are based on marginal properties; how the agent's actions change the metric. Consequently, a measure that provides perfect information can create highly distorted incentives, while a measure that gives excellent incentives can yield very little information. Because of this separation of the information and incentive properties of a metric, we simplify the metric choice problem into a choice of investment of resources into information or incentives. We show that the information about the agent type will dictate the principal's relative investment into information and incentives properties of the performance measure. More specifically, if the principal knows that the agent is of the 'World War II type' (using Kerr's examples), then it is optimal to invest all resources into information. Alternatively, knowledge of a 'Vietnam type' agent should result in a metric that equalizes the marginal value of information to the marginal value of incentives. We also, show that if the agent is of the 'Vietnam type' and the principal does not know this for sure, then the principal will believe it is closer to achieving its goals than is actually the case.

We develop the argument through the following steps. In the next section we discuss the application of principal-agent models and show how we adapt the general framework to fit the special case of military bureaucracies. Next we develop a formal model of military principal-agent interaction with imperfect measurement. We then illustrate the insights of the model through a brief contrast of the Second World War and the Vietnam War. We show that in the case of the Second World War, the agent had aligned incentives with the principal, the principal knew this to be the case, and it was well understood how the agent's actions impacted the wartime environment. This allowed for the simple choice of metrics that emphasized informational content. In the case of the Vietnam War, however, the principal was uncertain of the agent's incentives, and the impact of the agent's actions on the wartime environment was poorly understood. This resulted in poor metric selection and the inducement of pathological agent behavior.

In the final section we generate and explore policy implications for the ongoing conflict in Afghanistan. More specifically, we argue two things. First, that the incentives aspect of metrics has entirely been neglected; debate has uniformly focused on the improving the informational aspect of measurement. We show that this is a mistake. Second, the underlying cause of the metrics problem can be traced to a fundamental inability of the political-military leadership to articulate how military activity affects the Afghan operational environment. Until these issues are confronted, we predict a continued struggle for ‘good’ metrics in the conflict.

MILITARY AGENTS AND THEIR OPERATIONAL ENVIRONMENT

Before presenting the model, we describe the unique characteristics of the military agent and its relationship to its political principal. This is important, because the vast majority of principal-agent modeling has been done in the realm of profit-motivated firms and agents. We deviate from this standard principal-agent arrangement to better address the problem of professional soldiers. Bureaucracies are large, complex organizations who operate at the behest of a political principal. Due to these characteristics, formal principal-agent models have been increasingly used to model the interaction between political authorities and bureaucratic agents, which has contributed significantly to our understanding of these dynamics (see, for example, Calvert et al. 1989). Principal-agent models have rarely, however, been specifically crafted to deal with the nuances of the uniformed military services (one notable exception is Feaver 2005). Principal-agent models are especially attractive for modeling military bureaucracies in wartime environments in ways that have been hitherto neglected. In particular is the fact that political leaders' goals in wartime (such as ‘establish regional hegemony’ or ‘stabilization of a failed state’) often cannot be directly observed or contracted; rather, the principal must establish ‘observable benchmarks’ upon which progress towards the political endstate is noted.² We argue that these benchmark performance measures are analytically distinct from simple state-of-the-world information measures and we show the dangers of failing to parse out the relative effects of the two.

First, we define the military ‘agent’ as the uniformed services utilized to conduct operations in theater.³ For example, current US military operations in Afghanistan are

conducted under CENTCOM (Central Command - the regionally aligned combatant command). Conversely, we define the ‘principal’ as the military-political leadership apparatus. In the case of the United States, this would include the president and his senior advisors based in Washington DC.

Second, we depart from more standard economic and bureaucratic principal-agent models in our assumptions about military institutions. The standard hidden action models are predicated on the notion that agents prefer to ‘shirk than work’ - leveraging any asymmetric informational advantage to the most resources while putting in the minimum effort. This is what is commonly referred to as ‘information rent’ (Laffont and Martimort 2002: Ch 2). In the context of a professional military agent in wartime, however, this work-shirk paradigm is not necessarily appropriate.⁴ As Huntington's classic work on civil-military relations argues, the structure of a professional military institution is designed to prevent such venality, and members are also socialized away from economic motivation: “the employment of his expertise promiscuously for his own advantage would wreck the fabric of society... Clearly [the agent] does not act primarily from economic motivations” (1957: 14-15).⁵ For this reason, we do not assume that the principal must be concerned with a military agent’s choice of how much to shirk based on compensation.⁶ We, rather, assume that the military institutions guarantee that the agent utilizes any assets given to it, and that the principal is concerned simply with the nature and direction of the actions employed by the agent to execute the task.⁷ We show that understanding how performance measurement impacts agent incentives is fundamental even without the work-shirk paradigm.⁸ The problem we study is still driven by the unavoidable information asymmetry; the agent has ‘man-on-the-spot’ information in the field to make decisions about the direction of his actions, while the principal possesses more contextual information about the conflict. The principal is concerned with aiming agent behavior in the best direction, which is accomplished with performance metrics.⁹

There is a large literature on principal-agent relations beginning with Mitnick (1975).¹⁰ We build off the formal literature in employment relations starting with seminal contributions by Holmström (1979) and Grossman and Hart (1983). Our model is closely related to the conventional literature on agency problems of labor contracts with

imperfect performance measures. Baker (1992), for example, studies the effect of imperfect performance measures on the optimal proportion of performance pay and fixed salary. It is shown that the percentage of the contract that is performance pay is increasing in the correlation between performance measure and the objective function. These results are fundamental for the understanding of labor contracting in the private firm setting. We adapt Baker's model to the setting of military employment and focus on the trade-offs between informational content and incentives in performance metrics. The fundamental question Baker addresses is much less interesting in our context, since the optimal performance bonus is always arbitrarily small - such as medals, unit citations, promotions, or career assignments.¹¹

Finally, we clarify some terms before proceeding. First, we use the term 'metrics' inclusively to deal with task of assessment in wartime. This would include measures of 'performance,' 'progress,' and 'effect' as well as military intelligence collection efforts (see Daddis 2011: 5-16, Connable 2012: 2-3). Second, we refer to the 'operational environment' in which the conflict takes place. This would include the relevant aspects of the battlespace in which the war is being fought, and may include the geographical, technological, diplomatic, social context in which the war takes place (see Gray 1999: 23-44). This would accord with the meaning of the more general term 'state of the world' used in developing the formal model. We now proceed to the model.

THE MODEL

There are two actors: a principal and an agent. The principal has an objective function $v(a, \omega)$, which is not observable in the contracting time by the principal. It is a function of a , the actions of the agent, and ω , a vector of random variables that completely characterizes the state of the world. Denote by \mathbf{A} the set of possible actions and denote by Ω the set of all possible states of the world. The principal has probabilistic beliefs about the state of the world, which are specified by a probability measure f on Ω .

To help illustrate the range of choice in the set of actions \mathbf{A} , consider the following example: a particular action, a , has two dimensions, one specifies quantity of resources put into combat operations (such as kinetic strikes), the other specifies the resources put into noncombat operations (such as economic development projects). In

other words, military strategy involves choosing an action that involves multiple dimensions in an extremely complex operational environment. The range of action available to a commander across all such dimensions is huge, which underscores the need for the principal to effectively induce the desired behavior (on this issue of such latitude in military command, see Van Creveld 1985).

The principal's objective function v is such that there exists actions and states (a, ω) such that $v(a, \omega) = V$ and for all other (a', ω') , $v(a', \omega') < V$. The principal has a goal, which is $v(a, \omega) \geq V$ with probability one. If the goal is reached with probability one, then the conflict ends and there is no future cost. If the goal is not reached with probability one, then there is a positive expected future cost of conflict. We will address this cost more explicitly in Section 3.2.

The principal cannot observe the realization of ω , the agent's actions a , or the realized value of the function $v(a, \omega)$. The principal uses a metric p , chosen from the set of feasible metrics \mathcal{M} , to gain information about v and the goal. A metric p is a function from the space of actions and states of world to the real numbers.¹²

Before proceeding with the model, we provide a simple example to give a concrete illustration of the separation of information and incentives.

Example 1 (*Information versus Incentive*) Suppose that there are two possible states of the world: the enemy is 'weak' (denoted below as $\omega = 0$), and the enemy is 'strong' (denoted below as $\omega = 1$). The agent observes this and bases its choice on the opposition's strength. The actions available to the agent are 'direct attack' (denoted below as $a = 0$) and 'indirect attack' (denoted below as $a = 1$). The appropriate action is a direct attack if the opponent is weak ($a^0(0) = 0$), and an indirect attack if the opponent is strong ($a^0(1) = 1$). The principal has different values for different states and actions of the agent. These values are described below. Recall that the principal cannot observe the state, actions, or value. The goal of the principal is to get the value at or above 10 for sure; only at this point will the operations be concluded. Formally, the values $v(a, \omega)$ with action $a \in \{0,1\}$ and states $\omega \in \{0,1\}$ are

$$v(0,0) = 20, v(0,1) = 2, v(1,0) = 0, v(1,1) = 5.$$

First, we provide an example of a metric that has perfect incentives and provides no information to the principal. Metric 1 below is such a metric.

$$p^1(0,0) = 1, p^1(0,1) = 0, p^1(1,0) = 0, p^1(1,1) = 1$$

With metric 1, the agent maximizes the metric by picking the appropriate action for each state of the world. If the principal understands the incentive of the metric, then it will know that in state 0 the agent takes action 0, and in state 1 the agent takes action 1. In other words, the agent will choose a direct attack when the enemy is weak and an indirect attack when the enemy is strong; these are the appropriate responses within these two states of the world. The problem for the principal is that the two situations result in a metric value of 1. Thus, there is no way for the principal to know which state of the world has been realized; the principal does not know if the conflict is going exceptionally well (value 20) or not so well (value 5). Although the metric gives information in the weak sense that the principal knows one of these two cases must have occurred, the lack of information about the states of the world (and more importantly v) keeps the conflict continuing indefinitely, even in the case that the goal has actually been attained. Second, we provide an example of a metric that is fully informative about the state of the world, but gives the agent the wrong incentives in both states.

$$p^2(0,0) = 0, p^2(0,1) = 2, p^2(1,0) = 1, p^2(1,1) = 1$$

Based on metric 2, the agent will pick action 1 in state 0 and action 0 in state 1. These are the actions the principal would least like the agent to take. If the principal understands the incentives of the agent, however, then it will know that the agent will choose these actions. Since these outcomes have distinct metrics (action 0 in state 1 yields: 2, and action 1 in state 0 yields: 1), the principal can observe the metric and determine whether the state is 0 or 1. Thus, the metric is perfectly informative - the principal knows for sure the state and action in each state of the world. Although metric 2 provides perfect information, it will never incentivize the agent to perform the correct

action in state 0 and consequently the conflict will continue indefinitely. Third, we present a metric that has information and incentives properties less extreme than the first two examples.

$$p^3(0,0) = 2, p^3(0,1) = 1, p^3(1,0) = 0, p^3(1,1) = 1$$

In state 1, the agent might take either action and the principal is able to infer the state, but not the action. While in state 0, the metric gives the correct incentives and is informative to the principal. Thus, if state 0 occurs, then the agent will take action 0 and the principal will know that the value function is 20 and the operation has been successful.

The Agent's Problem

There are two types of agents: 0 and 1, with the arbitrary type denoted by the parameter θ , where $\theta \in \{0,1\}$. The agent knows its type and the objective function of the agent is dependent on the type. The principal uses linear incentive contracts based on the performance measure.¹³ The incentive contract for the agent for actions a in state ω is $bp(a, \omega)$, where b is the performance bonus.¹⁴ As we will see later, only one type of agent is motivated by the incentive contract. We assume that the agent observes the state of the world ω before the decision of actions $a \in \mathbf{A}$, where \mathbf{A} is a compact subset of \mathbb{R}_+^n for $n \geq 2$.¹⁵

The agent is given an endowment of resources r determined exogenously. We formalize the ability of military institutions to induce full effort utilization by assuming full resource utilization. That is, $a_1 + \dots + a_n = r$.

The agent type depends on the environment of the conflict. We will illustrate this in full when we cover the cases of World War II and the Vietnam War. The first type is an agent who maximizes the principal's objective function regardless of the metric. For this to be true an agent must be able to observe and understand v and have external incentives that make maximizing v the priority. On the other hand, if an agent cannot observe (or understand) v or is not externally incentivized to maximize v , then it can only maximize the performance bonus.

The distinct nature of military institutions greatly impacts the principal-agent interaction. We derive a simple, yet useful result regarding the optimal bonus b : the optimal bonus is arbitrarily small. Take any fixed endowment r and consider the optimal choice of the bonus b . Since the agent always utilizes full effort, for all $b > 0$, the agent will find it optimal to choose the same effort distribution. Since bonus payment is costly for the principal, the smallest positive bonus will give the principal the highest expected utility. Technically, this causes a problem of non-existence of a solution, since for all $b > 0$ there exists b' such that $b > b' > 0$. To rectify this closure issue, we take the bonus to be effectively costless to the principal.

At this point we drop the arbitrarily small bonus from the notation. We do this because we are assuming that: (i) the bonus is small enough that the agent of type 0 does not respond to it at all, and (ii) the bonus is still sufficient to direct the behavior of agent type 1 to just maximize the metric and get the largest bonus possible. We formalize these two types below.

The objective function of type $\theta = 0$ is the value function $v(a, \omega)$. The agent has fixed resources r and picks its actions to maximize v . Denote the optimal actions of type 0 by

$$a^0(\omega) \in \operatorname{argmax}_{a \in A} \{v(a, \omega) \mid \sum_{i=1}^n a_i = r\}.$$

The objective function of type $\theta = 1$ is the metric $p(a, \omega)$.¹⁶ The agent has fixed resources r and picks its actions to maximize p . Denote the optimal actions of type 1 by

$$a^1(p, \omega) \in \operatorname{argmax}_{a \in A} \{p(a, \omega) \mid \sum_{i=1}^n a_i = r\}.$$

We assume that all $p \in \mathcal{M}$ and v are continuous and the constraint set is compact, which guarantees the existence of a solution to both problems.¹⁷

The Principal's Problem

The principal does not necessarily know the realized type of the agent. It believes that the agent is of type 0 with probability $\mu \in [0,1]$ and of type 1 with probability $1 - \mu$.¹⁸

Given some metric p and realization of the metric ρ , we construct the principal's expected utility. First we must define some preliminaries.

A metric increase the *ex ante* expected payoff of the principal in two ways. First, if a metric is said to have better incentives, then that means the metric induces the agent of type 0 to take actions that result in a higher *ex ante* expected value for the principal. This increases the principal's expected payoff now as well as in the future. Second, if a metric is said to provide better information, then that means the *ex ante* expected future payoff of the principal is increased because (i) knowing more about the state of the world helps the principal use resources better in the future, reducing the cost of conflict and (ii) the increased knowledge of the conflict allows the principal to realize with greater clarity when its goal has been met.

For simplicity we reduce the set of possible metrics to functions characterized by two parameters (x, y) . We assume that the real numbers (x, y) capture all the relevant features about the metric. Thus a metric p is identified completely by its parameters (x, y) ; where x indicates the informational content of the metric and y indicates the incentive value of the metric. The possible values of (x, y) lie in the intervals $[0, X]$ and $[0, Y]$, where $x = X$ indicates a perfectly informative metric and $y = Y$ indicates a metric with ideal incentives.¹⁹

We make the following simplifying assumptions about the structure of the possible metric in order to provide parsimony to the analysis. Since all incentive features of the metric are captured by y , the agent's actions in type $\theta = 1$ are only impacted by the change in the parameter y . Formally, we can rewrite the function $a^1(y, \omega)$ instead of $a^1(p, \omega)$. Similarly, we impose that all information value comes through an increase in the parameter x .

Denote by $v^0(\omega) = v(a^0(\omega), \omega)$ and $v^1(y, \omega) = v(a^1(y, \omega), \omega)$. The principal's immediate expected utility of the metric is:

$$V^0 = \int_{\Omega} v^0(\omega) df,$$

$$V^1(y) = \int_{\Omega} v^1(y, \omega) df.$$

Since we have defined the parameter y as the incentive value of the metric, formally this means that an increase in y increases $V^1(y)$. For convenience we also assume that V^1 is twice-continuously differentiable and concave in y .

The second part of the principal's payoff is the future cost of the continued conflict. This cost is lowered by both the nature and precision of information. The cost is also impacted by incentives of the metric in the case that $\theta = 1$. We denote the two cost functions for parameters $\theta \in \{0,1\}$, $\Phi^0(x)$ and $\Phi^1(x, y)$. We assume that the future cost of the conflict is decreasing and convex in each information and incentive value of the metric. Formally, if $\Phi^0(x) > 0$ and $\Phi^1(x, y) > 0$, then all x and y are such that $\partial\Phi^0(x)/\partial x < 0$ and $\partial^2\Phi^0(x)/\partial x^2 \geq 0$, while $\partial\Phi^1(x, y)/\partial x < 0$ and $\partial^2\Phi^1(x, y)/\partial x^2 \geq 0$. Further, we assume that $\partial\Phi^1(x, y)/\partial y < 0$, $\partial^2\Phi^1(x, y)/\partial y^2 \geq 0$ and $\partial^2\Phi^1(x, y)/\partial x\partial y = 0$. These assumptions reflect the future value of information in the sense that more information helps the conflict end more quickly and with fewer losses. The second order properties of the cost function reflect that the marginal value of information is decreasing at higher information levels.

In addition we impose that the information parameter increases the cost of continued conflict in the same way in both types of agent. That is, $\partial\Phi^0(x)/\partial x = \partial\Phi^1(x, y)/\partial x$ for all x, y . This assumption is made to for the purpose of parsimony.

The cost to implement a metric is denoted by $C(x + y)$ and is increasing at an increasing rate: $C' > 0$ and $C'' > 0$. The cost assumption reflects the feature that the cost of implementing a better metric in either dimension is increasing and the marginal cost becomes larger as the aggregate level of information and incentives increase. The assumption also captures the increasing difficulty of actual implementation of a better metric.

Thus, the principal's expected payoff is

$$U(x, y) = \mu[V^0 - \Phi^0(x)] + (1 - \mu)[V^1(y) - \Phi^1(x, y)] - C(x + y).$$

The principal's maximization problem is

$$\max_{(x,y) \in [0,X] \times [0,Y]} U(x, y).$$

It is worth noting that we have constrained the principal to use a very restrictive class of pooling contracts as their metric. We have done so based on our beliefs that the institutional constraints of our application would not allow for more efficient sorting contracts. Further, we have assumed that no optimal pooling contract exists, i.e., v is not in the set of possible metrics.

Given some value for $\mu \in [0,1)$, the first order condition that defines all interior solutions of the optimal metric choice problem is $\partial U(x_\mu^*, y_\mu^*)/\partial x = 0$ and $\partial U(x_\mu^*, y_\mu^*)/\partial y = 0$ or

$$-\mu\Phi^{0'}(x_{0\mu}^*) - (1 - \mu) \frac{\partial\Phi^1(x_\mu^*, y_\mu^*)}{\partial x} - C'(x_\mu^* + y_\mu^*) = 0 \quad (1)$$

$$(1 - \mu) \left[V^{1'}(y_\mu^*) - \frac{\partial\Phi^1(x_\mu^*, y_\mu^*)}{\partial y} \right] - C'(x_\mu^* + y_\mu^*) = 0 \quad (2)$$

For any interior solution, the following intuitive relationship must hold

$$-\mu\Phi^{0'}(x_{0\mu}^*) - (1 - \mu) \frac{\partial\Phi^1(x_\mu^*, y_\mu^*)}{\partial x} = (1 - \mu) \left[V^{1'}(y_\mu^*) - \frac{\partial\Phi^1(x_\mu^*, y_\mu^*)}{\partial y} \right]$$

Marginal value of information *Marginal value of incentives*

The information value is based on how much the realization of the metric allows the principal to infer about the state of the world and consequently the value function. More knowledge about the state of the world allows the principal to make better decisions in the future that lower the future cost of conflict as well as recognize the goal has been attained sooner. The value of better incentives is based on inducing a type 1 agent to take actions that are better for the principal.

In what follows, the model is used as the basis of understanding the special cases that correspond to the Second World War and the Vietnam War.

VARIATION ACROSS WARTIME ENVIRONMENTS

We illustrate the model's insights by comparing heterogeneous wartime cases. The United States' participation in the Second World War and the Vietnam War provide

significant variation on the parameters of interest. In doing so, we show that each of these historical conflicts accord with a special case of the model. The Second World War accords with the case in which the agent's incentives align with the objective function of the principal. Further, it was a wartime environment in which the agent could easily apprehend the objective function and how its efforts contributed to progress towards the principal's goal. Finally, the principal knew the agent's type - it understood the nature of the agent it was commanding. In Vietnam, these things were not the case: the agent did not understand the wartime environment and it did not apprehend how its efforts would contribute to achieving the principal's underlying goals. Further, the model stipulates that both the agent's incentives must align with the principal's objective function and that the principal knows the agent's type. These necessary conditions were not met; ergo, metric selection was problematic. We now demonstrate the model through these illustrative cases.

The Second World War

We argue, in the case material that follows, that World War II fits the parameter values $\theta = 0$ (the agent's incentives are aligned with those of the principal), and $\mu = 1$ (the principal knows for sure the agent's type). If $\theta = 0$ and $\mu = 1$, then the marginal utility through incentives is zero. This is because, regardless of metric, the agent always maximizes the objective function in every state of the world. In this case, the optimal metric selection reduces to maximizing the information value alone. The principal can focus solely on attaining as much information as possible to understand when the goal has been met and thereby reduce future cost.

Since it is optimal to set $y_1^* = 0$, the first order conditions for the optimal metric choice reduce to

$$-\frac{\partial \Phi^0(x_1^*)}{\partial x} = C'(x_1^*).$$

This is a relatively straightforward problem; choose a metric that balances benefits of more information with the cost of implementation.

Second World War Empirical Case

To apply the model to the case of the Second World War requires exploring the values of the three parameters. The first is θ ; this is the agent type. In this case, the agent's incentives were aligned with the principal's objective function (v). Second, is in regards to the objective function itself, which, in this case, could easily be observed and understood by the agent. The third is μ , the principal's belief regarding the agent's type; in this case, that the principal had knowledge that agent's type was zero. In sum, the principal in the Second World War was not bedeviled by malign agent incentives, and could concentrate its metrics efforts on gathering information.

A crucial aspect of the Second World War was that the agent was incentive-aligned to the principal's conception of victory. By the end of the war, the size of the United States Army had risen to 8,266,373, from a pre-war size of just 269,023 (for more detailed discussion on this process of expansion see Koistinen 2004, and Newland and Chun 2011). The vast majority of these soldiers were draftees, who had been conscripted for the duration of the conflict (Chambers 1999: 181). This conscription mechanism aligned the principal's goal for victory with the agents' goal of resuming civilian life. To reiterate Kerr's encapsulation of this incentive for the soldiers: "What did the GI in World War II want? To go home. And when did he get to go home? When the war was won!" (Kerr 1975: 771).

A second component of the agent's type is that the agent understands the nature of the principal's value function and how its effort contributes to progress towards that goal. The early twentieth century represented the marriage of Napoleonic military strategy, nationalism, and the industrial revolution. Success in such warfare was predicated on destroying enemy forces, crippling the enemy society, and occupying its territory; President Roosevelt justified this wartime strategy to the US Congress: "We wage total war because our very existence is threatened...[it] is a grim reality...In total war we are all soldiers" (quoted in Heuser 2010: 194). Given this operational environment, the war was a relatively simple (though by no means easy) affair. The US Army built a doctrine in the inter-war period that embodied this understanding and provided a blueprint for agent activity:

World War I gave no promise that victory in modern war could grow from anything but the application of superior resources, not in dazzling maneuver... but in hard fighting. In the army's professional school system [throughout the inter-war period]... the war was fought and refought again and again... and the emphasis always was on the intractability of modern strategic problems to any solution save that of overwhelming power (Weigley 1986: 269).

Given this clarity, the principal had a high degree of certainty of the agent type - and was able to resource and reward the agent effectively. The exemplars of this type of war were General George Patton and General Curtis LeMay. Patton let his subordinates know clearly what was expected of them in a famous speech to his Third Army troops in May of 1944: "There is only one tactical principle which is not subject to change. It is to use the means at hand to inflict the maximum amount of wound, death, and destruction on the enemy in the minimum amount of time." This was normal fare for Patton, who wrote in his diary that year: "Made a talk [today]. As in all my talks I stressed fighting and killing" (quoted in Overy 1995: 173). Patton was as good as his word, becoming one of the most feared battlefield commanders of the war: "The statistical imbalance... was staggering... In total casualties - dead, wounded, and captured - the Third Army [under Patton] caused the enemy ten times the losses that it suffered - by far the greatest ratio of damage inflicted versus losses incurred in the entire Anglo-American force" (Hanson 1990: 303). Similarly, LeMay was considered a successful operational leader in the war; he was directly responsible for fire bombing 63 Japanese cities, killing a half-million Japanese civilians, and de-housing another 8 million. In short, "[f]or Lemay, demolishing everything was how you win a war" (Kaplan 1983: 43). Similar to Patton, LeMay was judged a successful commander for utilizing assets for maximum destruction (he was later Air Force Chief of Staff and was the youngest four-star general in modern history). These agent activities, then, could be linked directly back to the war's operational benchmarks that had been established months before the Pearl Harbor attack: control of the seas, operational air superiority, disruption of enemy industry, and ultimately the destruction of enemy military forces (see Kirkpatrick 1992: 63-77).

In sum, the United States war effort from 1941-1945 exemplified one case of the model. The agent's incentives aligned with those of the principal, the agent could grasp the nature of the true value function, and the principal was aware of the agent type. Taken together, these parameter values allowed the principal to focus his measurement efforts towards the gathering of information.²⁰ Such an attractive wartime environment has not always been the case in US military efforts. We now proceed to a conflict that had very different characteristics.

The Vietnam War

We argue in what follows that the Vietnam War fits a second case of our model: that $\theta = 1$ and $\mu > 0$. That is, the agent is of type 1 and the principal is not sure of the agent's type. In this case, the principal is in a position where it must consider the trade-off between information and incentives. If the principal does not know that the agent's parameter is 1, then this leads to the principal choosing a metric that overvalues information, compared to the case when the principal knows that $\theta = 1$. We show this formally in the following proposition. In what follows we assume that there is an interior solution for all $\mu \in (0,1]$.

Proposition 1 *If the principal's beliefs are $\mu > 0$ instead of $\mu = 0$ when $\theta = 1$, then $x_\mu^* > x_0^*$ and $y_\mu^* < y_0^*$.*

Proof of Proposition 1. The argument for why the principal over-invests in information and under-invests in incentives is based on using the first order conditions (1) and (2). Take the optimal metric at $\mu = 0$, (x_0^*, y_0^*) . Now let us consider $\mu > 0$ and the choice of an optimal metric. We will show that the optimal metric for $\mu > 0$ is such that $x_\mu^* > x_0^*$ and $y_\mu^* < y_0^*$. Let us begin with considering metric (x_0^*, y_0^*) for some $\mu > 0$. We will construct a sequences of parameters and define the first element of this sequence by $x^1 = x_0^*$ and $y^1 = y_0^*$. The marginal value from incentives is decreasing in μ . Therefore, at x_0^* it must be that the incentive parameter is some $y^2 < y_0^*$ to satisfy (2) with $\mu > 0$. Based on the strict concavity of the objective function in y , such a y^2 is unique. At y^2 , the marginal cost is lower at x_0^* and consequently the marginal value of information must be

decreased to satisfy (1). Since the marginal value of information is non-increasing in x , x must be increased to $x^2 > x_0^*$ to equalize (1). Based on the strict concavity of the objective function in x , such a x^2 is unique. We continue with an iteration of the same process; $y^3 < y^2$ must be picked to satisfy (2) at x^2 . At y^3 , the marginal cost is lower and the value of x must be increased to x^3 to equalize (1). These directions of change continue with each iterative step, and we use them to construct two sequences $\{x^k\}$ and $\{y^k\}$. Each sequence is strictly monotonic and with each element of the sequence contained in a compact space ($[0, X]$ and $[0, Y]$, respectively). Any strictly monotonic sequence in a compact space must converge to some point in that space. Further, the sequences must converge to the unique equilibrium (x_μ^*, y_μ^*) . based on the continuity of the first order conditions. Since x^k is a strictly increasing sequence and y^k is a strictly decreasing sequence, it must be that $x_\mu^* = \lim_{k \rightarrow \infty} x^k > x^1 = x_0^*$ and $y_\mu^* = \lim_{k \rightarrow \infty} y^k < y^1 = y_0^*$. ■

Further, not knowing that the agent is type 1 creates an informational bias that skews the principal towards believing the conflict is going better than is actually the case. In Proposition 2, we prove that if the principal does not know for sure the agent is type $\theta = 1$, than the principal's beliefs about the conflict will be skewed towards over optimism. We make this argument with the metric choice fixed. If we also consider that the metric is skewed towards information (from Proposition 1), then the over optimism problem is further exacerbated.

The proposition is based on the following reasoning. The principal receives information from the realized value of the metric, but cannot observe the underlying state of the world and actions. The principal thinks there is some probability that the agent is type 0 and in that case, for such a realization of the metric, the actual level of success would be higher. Since the agent is actually type 1, the principal's expected value of actual success is higher than reality. This is what drives the optimistic bias.

In order to show this result we add more structure to our set of possible metrics. The following properties are used to prove the proposition. First, a metric p is *varied* on the state space Ω if, for all $\omega \in \Omega$ there exists $\omega' \in \Omega$ such that $p(a^0(\omega), \omega) = p(a^1(\omega'), \omega')$. This means the set of states of the world is varied enough

that there is a state such that a metric can come up with any value given an agent of type 0 or 1. Second, a metric p is *state consistent* with v if for all states $\omega, \omega' \in \Omega$ and actions a , $p(a, \omega') > p(a, \omega)$ if and only if $v(a, \omega') > v(a, \omega)$. This means that, fixing the actions of the agent, any state that improves the metric similarly improves the principal's objective.

Proposition 2 *Suppose that the metric is varied and state consistent with v and does not have perfect incentives (i.e., $y < Y$). If $\theta = 1$ and the principal believes that there is some probability that $\theta = 0$, then the principal expectation of v will be higher than if it knew $\theta = 1$ for sure.*

Proof of Proposition 2. Denote ρ as the realized value of the metric. Based on the fact that p is varied on Ω there exists the non-empty sets of states $\Omega^0(\rho)$ and $\Omega^1(\rho)$ such that $p(a^0(\omega), \omega) = p(a^1(y, \varphi), \varphi) = \rho$ for all $\omega \in \Omega^0(\rho)$ and $\varphi \in \Omega^1(\rho)$. Based on the fact that $a^0(\omega)$ is defined as a maximizer of v at each ω , $v(a^0(\omega), \omega) > v(a^0(\varphi), \varphi)$ for all $\omega \in \Omega^0(\rho)$ and $\varphi \in \Omega^1(\rho)$. Further, we know by the optimality of a^0 that $v^0(\varphi) > v^1(y, \varphi)$ for $y < Y$. Putting these together we have that

$$\int_{\Omega^0(\rho)} v(a^0(\omega), \omega) df > \int_{\Omega^1(\rho)} v^0(\omega) df > \int_{\Omega^1(\rho)} v^1(y, \omega) df \quad (3)$$

Note that

$$E[V^0 | \rho, \theta = 0] = \int_{\Omega^0(\rho)} v(a^0(\omega), \omega) df,$$

$$E[V^1(y) | \rho, \theta = 1] = \int_{\Omega^1(\rho)} v^1(y, \omega) df,$$

and (3) implies that $E[V^0 | \rho, \theta = 0] > E[V^1(y) | \rho, \theta = 1]$. Based on the fact that we have assumed that the observation of ρ does not impact μ , we can use these to compare the

expected values when $\mu = 0$ and $\mu > 0$ and complete our proof. Note that

$$\begin{aligned} E[V^\theta(y)|\rho, \mu = 0] &= E[V^1(y)|\rho, \theta = 1] \\ &< \mu E[V^0|\rho, \theta = 0] + (1 - \mu)E[V^1(y)|\rho, \theta = 1] \\ &= E[V^\theta(y)|\rho, \mu > 0]. \end{aligned}$$

■

Now we apply these results to understating the conflict in Vietnam.

Vietnam War Empirical Case

The case of Second World War can be contrasted with that of the Vietnam War. In this vastly different operational environment, US military leaders attempted to defeat a political insurgency and establish a viable South Vietnamese government through the operational benchmark of ‘winning the hearts and minds’ of the population. The core tactical metric, however, was the use of ‘body counts’ to attrit enemy forces to the degree that they could no longer replace their losses. The pathology engendered by this metric choice, however, was that it incentivized large-scale killing and destruction, which worked against the goal of building a viable political regime in the South. It is important to emphasize that killing is understood to be a defining aspect of any war; in the case of Vietnam, however, as the principal was unaware of the agent's type, the pursuit of the metric for its own sake reduced the informational content provided by the reported body counts. In other words, while counting enemy dead provided important information in the Second World War context, it failed to do so in Vietnam. Further, our model helps to explain why the military in Vietnam over-produced violence, and why the political leadership was unduly optimistic given the assessments from field.

Once again, it is important to stipulate that the Vietnam military challenge contained multiple dimensions that needed to be addressed: military, political, economic, and social. Any optimal action choice, a , would need to be balanced appropriately across these dimensions. Carland summarizes the problem thusly: “when military victories were won, their significance lay in the degree to which they advanced and supported South Vietnam’s pacification/nation-building effort... if they failed to integrate the ‘fighting’ war with the ‘other’ war they would not succeed” (2004: 554). In other words, regime stabilization was based on several non-military dimensions that spread far beyond combat

efforts and in which killing should have only played a supporting role (Rosenau 2005; Jones 2012). The ultimate goal was to establish a viable, self-sustaining polity that would provide a bulwark against further communist ‘dominoes’ falling across South East Asia.²¹

The US military agent in the Vietnam War, however, did not have incentives that aligned with the principal’s goal of stabilization. Instead, the agent was driven toward pursuing the performance metric of killing. This was true for at least two reasons. First, soldiers served one-year tours of duty in combat (the ‘rotation system’). The incentive for these soldiers was to simply survive this period and go home, regardless of whether progress was made towards victory: “the rotation system reinforced an individualistic perspective that was essentially self-concerned. The end of the war was marked by the individual's rotation date and not by the war's eventual outcome--whether victory, defeat, or negotiated stalemate” (Moskos 1975: 31). The second set of incentives was for career officers and noncommissioned officers. In fact, Moskos argues that this rotation system drove this novel bifurcation within the Army organization: “where army internal cleavages had formerly derived from the basic distinction between enlisted men and commissioned officers, the emergent distinction became that between single-term soldiers - whether officer or enlisted - and career soldiers - whether officer and enlisted” (1975: 32). For these career-oriented ‘lifers’ the incentive in Vietnam was to maximize performance metrics for the purpose of earning citations and promotions during their rotation. In sum, officers and units were driven to maximize the performance metric while the drafted personnel were incentivized to simply survive. As we will show below, these twin dynamics created an agent that was driven to overproduce violence and casualties, rather build a sustainable South Vietnamese regime.

Even if agent incentives had aligned with the objective function of the principal, it would not have been able to pursue that goal very effectively. It was simply the case that the US military could not determine how to best use its resources in Vietnam.

US policy makers had outlined national objectives, such as South Vietnamese independence and territorial integrity, countering Communist influence and pressure, and controlling insurgent elements. Clear objectives for the use of

military force, though, never accompanied these general goals, leaving the armed forces searching for linkages between strategy and policy (Daddis 2011: 47).

In the absence of establishing this linkage between force and policy the army defaulted back to its traditional way of war: “When General Westmoreland was asked at a press conference what the answer to insurgency was, his reply was one word: ‘Firepower’” (Krepinevich 1986: 197).

The Army's default to conventional warfighting techniques matched with a performance measure based on killing: the infamous ‘body count.’ Secretary of Defense Robert McNamara was the originator of this metric:

I insisted we try to measure progress... I was convinced that, while we might not be able to track something as unambiguous as a frontline, we could find variables that would indicate our success or failure... Critics point to use of the body count as an example of my obsession with numbers... Obviously, there are things you cannot quantify... [b]ut things you can count, you ought to count. Loss of life is one when you are fighting a war of attrition. We tried to use body counts as a measurement to help us figure out what we should be doing in Vietnam to win the war while putting our troops at the least risk (McNamara 1995: 237-238).

Given this imposed metric of performance, the agent pursued it accordingly. Appy writes that these “death tallies were constantly monitored and updated. In rear areas, command posts listed ‘box scores’ on large chalkboards... Indeed, killing was the central focus of American policy” (1993: 144). In turn, as Shelby Stanton writes, units and officers were “rewarded by promotions, medals, and time off from field duty. For example, General Westmoreland had issued a special commendation to the 11th Infantry Brigade based on its claim of 128 killed at My Lai [these victims turned out to be civilians, in what was later deemed the ‘My Lai Massacre’]” (quoted in Gartner 1997: 128-129). Such gross levels of violence was inimical to the ultimate goal of ‘winning hearts and minds’ of the South Vietnamese people, yet US forces in Vietnam were incentivized to engage in such indiscriminate killing: “[T]he Army maintained that it closely observed very restrictive

rules of engagement (ROE) throughout the war... Yet, by placing the body count above population security in its list of priorities, the Army provided the incentive for its commanders to shoot first and worry about the hearts and minds later” (Krepinevich 1986: 198-199, emphasis added). In other words, the support of the civilian population was sought, but performance metrics incentivized agent behavior that was very divergent from the path towards the principal’s goal.

It is unclear whether the principal knew of the nature of its agent in Vietnam. For one thing, the principal had very poor understanding of what the agent was doing in South East Asia: “lack of expertise hampered the ability of the administration to hold the Army's feet to the fire over counterinsurgency; thus, the Army could give lip service to requirements placed on it by the administration or ignore them entirely” (Krepinevich 1986: 33). This resulted in fundamental lack of understanding for the principal as to how the agent's effort was tracking with actual progress towards its goals. Secretary of Defense McNamara admitted that it

was not the valor of American soldiers in Vietnam that was ever in dispute but how they should operate in the field. This issue became the focus of considerable disagreement between Westy [Army general Westmoreland] and the marines... Although deeply divided, the military never fully debated their differences in strategic approach, or discussed them with me in any detail. As secretary of defense, I should have forced them to (McNamara 1995: 243, emphasis added).

As the model predicts, the principal's uncertainty over agent type resulted in systematic over-optimism concerning the war's progress. This was highlighted most clearly by the infamous ‘light at the end of the tunnel’ pronouncement, made by Westmoreland in November of 1967. In media interviews and at the National Press Club during a trip to Washington, Westmoreland exuded confidence and, in turn, briefly buoyed public perception regarding the war: “With a definite end of the war in sight, the American public caught some of the optimism... even the popularity of President Johnson, which had been on a long downward spiral, recovered 10 ten points in one month... after General Westmoreland’s optimistic trip” (Blood 2005: 41). This optimism

was not just reserved for the public: “A ‘we are winning’ consensus pretty much permeated the Saigon-Washington command circuit...” (Ford 1998). This sense of confidence in the war’s progress was based on the body counts. This stemmed from the argument that as soon as enemy deaths outpaced the ability of the enemy to recruit new soldiers, the war was essentially won (it had reached the ‘crossover point’), and official statistics showed that the crossover point had indeed been reached in December of 1966 (Blood 2005: 29).²² This fits with our model: the principal did not know that the agent was pursuing the metric (body counts) rather than the objective function (a stable South Vietnamese regime), and as a result was overly optimistic regarding the war’s progress as enemy deaths mounted.²³

In sum, the Vietnam War exemplifies the problematic case of the model. The agent’s incentives differ from those of the principal and the principal does not know the agent’s type. In this case the agent pursues its performance measure, regardless of whether this activity contributes to the principal’s goal. In fact, in the case of Vietnam, the agent’s performance measure (body counts) worked against the principal’s goal of stabilizing South Vietnam. Two decades after the war, Secretary McNamara recognized this pathology: “Westy’s attrition strategy relied heavily on firepower... It often proved difficult to distinguish combatants from noncombatants. Fighting produced more and more civilian casualties... [and this]... undermined, in an unintended but profound way, the pacification program designed to... win the ‘hearts and minds’ of the South Vietnamese people” (McNamara 1995: 243).

CONCLUSIONS AND RECOMMENDATIONS

We have explored the dangers of neglecting to select metrics in wartime carefully. In particular, we developed a model to understand how imperfect measures of success may have deleterious externalities by creating unintended incentive structures for agents within the organization. Through a principal-agent analysis specifically tailored for application to military organizations in wartime, we have shown that the informational properties of the measurements are based on how the measure differs from operational success while the incentive properties of the measurement are based on differences in the marginal sensitivity of both the measure and operational success. We have shown that if

the agent's incentives align with the principal's goal, the principal knows the agent's type, and the agent understands how his actions affect the value function, then effective measurement is possible. Further, we then provided a framework to show the trade-off between information and incentives. Finally, we have shown that an under-appreciation of the incentive properties of measurement will lead to systematic positive bias of information.

We explored how two theoretical cases of the model accord with two historical cases from US foreign policy. In the case of the Second World War, the incentives of the agent aligned with that of the principal. Further, the agent observed the value function and grasped how his efforts contributed to its pursuit. Neither of these held true in the case of the Vietnam War. In the absence of aligned incentives or the comprehension of how to progress toward the principal's underlying goal, the agent reverted to maximizing the performance metric. In Vietnam, the overriding performance metric was the 'body count'; this led to a pathological over-production of violence, which actually worked against the principal's goal of a pacified and stable South Vietnamese regime.

Our analysis suggests that assessment efforts in the current conflict in Afghanistan exhibit many similarities to those of the Vietnam War. More specifically, that the military agent may be unable to grasp or observe progress towards the principal's goal of establishing a pacified and stable regime. Further, in the case of Afghanistan, the principal seems as unsure of the incentives of the agent. Servicemen and women - many of whom have served multiple tours in both Iraq and Afghanistan - are exhausted and have little reason to be optimistic about the long-term effects of their efforts in Afghanistan. It is reasonable to question whether they have reverted to their professional interests, and whether the principal would be aware of this development (see Bleigh et al. 2011).

Our model suggests that if the principal is uncertain as to the agent's type, he will most likely be led to incorrect conclusions regarding the war effort. He may seek detailed information about the war, but due to the disconnect between agent activity and the conflict's true progress, he will not understand the war. Consider the case of the Vietnam War: "Left with insufficient foundational knowledge of counterinsurgencies and vague strategic objectives, MACV [US Military Assistance Command, Vietnam] embraced

Secretary of Defense McNamara's advice that everything that was measurable should in fact be measured... Consequently MACV - and much of the DoD - went about measuring everything and, in a real sense, measured nothing” (Daddis 2011: 10). This sounds comparable to the current conflict in Afghanistan; as Kapstein argues, allied forces in Afghanistan are similarly attempting to build a “comprehensive data-set... Unfortunately, these metrics provide little more than a hodgepodge of trends, data, and ‘atmospherics,’ and [yet] its unclear how they relate to the war effort. In fact, this grab-bag of evidence suggests only one thing: that coalition forces still don't know how to measure their progress” (Kapstein 2011, emphasis in original). Our work further predicts that this continued effort will likely result in an upwardly biased assessment of the progress of the conflict - the principal may myopically believe that he is seeing a ‘light at the end of the tunnel’ or that the war has ‘turned a corner’ - when it, in fact, has not. Further research would be necessary to assess the degree to which the pathologies we have highlighted here do indeed exist in the Afghanistan conflict (for a critical overview of such efforts, see Connable 2012).

Finally, it is important to re-emphasize that our model is not restricted to these particular conflicts, but to wartime assessment in general. We recommend that the political leadership reorient its efforts to assess the military progress in all conflicts, taking into account incentives rather than relying on ever-greater levels of information. More specifically, we recommend searching for measures that are sensitive to action-choices of the military agent in as similar a way as possible to the ultimate political goal across as many states of the world as possible. The goal is to incentivize the agent to pursue the best interest of the principal, through pursuit of the performance metric, rather than to produce divergent - and perhaps counterproductive - behavior.

Acknowledgement

*We are heavily indebted to Aric Shafran for his many helpful comments on an earlier version of this paper. Rewadee Anujapad and Tricia Sullivan provided comments on clarity and presentation. Finally, we would like to thank the editor Daniel Arce and the two anonymous referees for their helpful suggestions that dramatically improved this article.

END NOTES

1. There is a significant literature on principal-agent relations with imperfect performance measure. Gibbons (1998) provides an excellent overview of some of this literature. Our model is most similar to the model of Baker (1992), which is used to examine the optimal balance of bonus and fixed salary compensation. In Section 2 we discuss the military-specific assumptions that separate our work from the literature pertaining to civilian labor contracts.
2. In the strategic studies literature, this is referred to as relating the goals of 'military strategy' (effective use of force) to the higher goals of 'grand strategy' (the country's overall foreign policy). On these points see Millett, Murray, and Watman (1986).
3. We find that treating the military as a single agent is a useful abstraction to provide focus and clarity. Of course, the 'military agent' is actually an immense organization teeming with internal agency problems. The nature of these internal organizational problems is outside the scope this treatment.
4. Feaver retains the term 'shirk' in his analysis, but redefines it to accord very closely with our approach - in which the principal is concerned with the nature of the agent's work rather than being 'lazy' as in the common usage of the term (2005: 58-68).

5. For a deeper analysis of this point, see the central argument of North, Wallis, and Weingast (2009). Their definition of the 'open access order' society hinges on such professionalized agents that monopolize organized violence.
6. As we will show in the case of the Vietnam War, soldiers served one-year tours of duty in combat (the 'rotation system'). The incentive for these soldiers was to simply survive this period and go home, regardless of whether progress was made towards victory (Kerr 1975). This behavior could be construed as 'shirking,' but for the purpose of survival, not personal enrichment. We assume away this phenomenon in the present study by unifying agent activity to the actual choice of tactics and operations employed by the unified military structure.
7. Permitting low effort 'shirking' would inherently make the incentive component of the metric more important. Our treatment shows the importance of incentives in spite of there being no concerns with inducing high effort.
8. The analysis could be done without assuming that military institutions remove the possibility of shirking. In contrast to our results, such a model would be predicated on significant compensation to the agent based on performance, something we do not see in modern, professionalized military organizations. Otherwise, a more standard hidden action model would imply similar results.
9. In more standard principal agent models the principal is the residual claimant - keeping the proceeds not given to the agent. Our model does not fit into this paradigm as our principal is assumed to have a fixed budget and have the sole goal of military success.

10. Mitnick (1975) provides the first formal treatment of the agency problem; the principal uses resources to direct behavior of a self-interested agent in the presence of information asymmetry.

11. Another relevant literature examines the problem of agency where the agent both gathers information and takes action based on the information. In the current manuscript we do not deal with issues of information acquisition by the agent, but instead focus on information gathering by the principal for later strategic decisions. It is very easy to see how the two issues could become convoluted. We treat information as verifiable in order to avoid the problem of delegated expertise and keep our treatment parsimonious. The literature dealing with information gathering in agency includes: Demski and Sappington (1987), Barron and Waddell (2003), Feess and Walzl (2004), Gromb and Martimort (2007), Malcomson (2009), and Inderst and Ottaviani (2009).

12. It is common for military to make a distinction between ‘measurements of performance’ (MOPs) and ‘measurements of effects’ (MOEs) (Center for Army Lessons Learned 2010). The model we develop is general enough to accommodate both MOPs and MOEs: An MOP would only be a measure of how the agent acted (i.e., $p(a)=p(a,\omega)$ for all ω) while the an MOE would be a measure that only attempts to approximate v . It is useful to see that the notion of MOP and MOE are not separable: by definition, agent actions impact operational success.

13. In our model assuming that the principal uses a linear performance bonus is without loss of generality, since such a bonus will be optimal.

14. We abstract away from modeling dynamic incentives of the agent, but the bonus b is intended to capture the future benefit to the agent through potential promotion and

advancement within the military organization, which is likely to provide the agent more expected utility in the future. We do not model this explicitly to avoid introducing unnecessary complexity to our model.

15. It is not important that the agent knows everything about the state of the world, just that the agent is in some way more informed about the state than the principal.

16. The agent maximizes the bonus $bp(a,\omega)$. Since b is a positive constant, this is the same as just maximizing the metric $p(a,\omega)$.

17. Since we have assumed that the military institutions guarantee $a_1+\dots+a_n=r$, we could include a cost of actions $C(\sum_{i=1}^na_i)$ with no ill effects.

18. Throughout the analysis we take μ as an exogenous parameter. This is done to avoid dealing with metric choice revealing the type of the agent; including this issue would convolute our main results.

19. Assuming that the key features of the metric can be summarized by two real numbers avoids the complexity of a choice problem over arbitrary function spaces. Similar results could be derived with weaker assumptions, but the results would lose the parsimony of the present treatment.

20. For a detailed case study of the efforts to measure effects in the allied strategic bombing campaign in Europe, for example, see Ehlers (2009).

21. Whether this strategy was appropriate or not is still the subject of debate. For a recent critique of this strategic approach, see Andrade (2008). For a brief overview of the vast historiography of the Vietnam War, see Hess (1994). 22. For further discussion of estimating the enemy order of battle, see Wirtz (1991).

23. Beyond the problem of incentivizing the over-production of actual killing, there was also the problem of inflated numbers of kills being reported: “It is generally assumed that the body count was exaggerated: there was considerable incentive for US and South Vietnamese officers to err on the high side or even to fabricate wildly to impress superiors” (Mueller 1980: 503, see also Travis 1990). This self-reporting problem lies outside the current model, but bears further theoretical development.

REFERENCES

- Andrade, D. (2008) Westmoreland was right: Learning the wrong lessons from the Vietnam War. *Small Wars & Insurgencies* **19** (2), 145-181.
- Appy, C.G. (1993) *Working-Class War: American Combat Soldiers in Vietnam*. Chapel Hill: University of North Carolina Press.
- Baker, G.P. (1992) Incentive contracts and performance measurement. *Journal of Political Economy* **100** (3), 598-614.
- Barron, J. M. and Waddell, G. R. (2003) Executive rank, pay and project selection. *Journal of Financial Economics* **67** (2), 305-349.
- Bleigh, J., Hufnagel, J. and Curt S. (2011) Institutional challenges to developing metrics of success in irregular warfare. Unpublished thesis. Monterey, CA: Department of Defense Analysis, Naval Postgraduate School.
- Blood, J. (2005) *The Tet Effect: Intelligence and the Public Perception of War*. London: Routledge.
- Calvert, R.L., McCubbins, M.D. and Weingast, B.R. (1989) A theory of political control and agency discretion. *American Journal of Political Science* **33** (3), 588-611.
- Campbell, J., O'Hanlon, M. and Shapiro, J. (2009) How to measure the war. *Policy Review* **157** (Oct/Nov), 15-30.
- Carland, J.M. (2004) Winning the Vietnam War: Westmoreland's approach in two documents. *Journal of Military History* **68** (2), 553-574.

Center for Army Lessons Learned. (2010) *Assessment and Measures of Effectiveness in Stability Ops*. Ft Leavenworth, KS: Combined Arms Center.

Chambers, J.W. (1999) Conscription. In the *Oxford Guide to American Military History*, edited by J.W. Chambers. New York: Oxford University Press, 180-182.

Connable, B. (2012) *Embracing the Fog of War: Assessment and Metrics in Counterinsurgency*. Santa Monica: RAND.

Daddis, G.A. (2011) *No Sure Victory: Measuring US Army Effectiveness and Progress in the Vietnam War*. New York: Oxford University Press.

Demski, J. S. and Sappington, D. E. M. (1987) Delegated expertise. *Journal of Accounting Research* **25** (1), 68-89.

Ehlers, R.S. (2009) *Targeting the Third Reich: Air Intelligence and the Allied Bombing Campaigns*. Lawrence, KS: University Press of Kansas.

Feaver, P.D. (2005) *Armed Servants: Agency, Oversight, and Civil Military Relations*. Cambridge: Harvard University Press.

Feess, E. and Walzl, M. (2004) Delegated expertise---when are good projects bad news? *Economics Letters* **82** (1), 77-82.

Flynn, M.T., Pottinger, M. and Batchelor, P.D. (2010) *Fixing Intel: A Blueprint for Making Intelligence Relevant in Afghanistan*. Washington DC: Center for a New American Security.

Ford, H.P. (1998) *CIA and the Vietnam Policymakers: Three Episodes 1962-1968*. Accessed at: <https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monographs/index.html>

Gartner, S.S. (1997) *Strategic Assessment in War*. New Haven: Yale University Press.

Gibbons, R. (1998) Incentives in organization. *Journal of Economic Perspectives*. **12** (4), 114-132.

Gordon, M.R. and Trainor, B.E. (2013) *The Endgame: The Inside Story of the Struggle for Iraq, From George W. Bush to Barack Obama*. New York: Vintage.

Gray, C.S. (1999) *Modern Strategy*. New York: Oxford University Press.

Gromb, D. and Martimort, D. (2007) Collusion and the organization of delegated expertise. *Journal of Economic Theory* **137** (1), 271-299.

Grossman, S. J. and Hart, O. D. (1983) An analysis of the principal-agent problem. *Econometrica* **51** (1), 7-45.

Hanson, V.D. (1990) *The Soul of Battle*. New York: Free Press.

Hess, G.R. (1994) The unending debate: Historians and the Vietnam War. *Diplomatic History* **18** (2), 239-264.

Heuser, B. (2010) *The Evolution of Strategy: Thinking War from Antiquity to the Present*. New York: Cambridge University Press.

Holmström, B. (1979) Moral hazard and observability. *Bell Journal of Economics* **10** (1), 74-91.

Huntington, S.P. (1957) *The Soldier and the State: The Theory and Politics of Civil-Military Relations*. Cambridge: Belknap Press.

Inderst, R. and Ottaviani, M. (2009) Misselling through agents. *American Economic Review* **99** (3), 883-908.

Jones, F.L. (2012) The guerilla warfare problem: Revolutionary war and the Kennedy administration response, 1961-63. In *US Army War College Guide to National Security Strategy, Vol.2*, edited by J.B. Bartholomees. Carlisle Barracks, PA: Strategic Studies Institute.

Kaplan, F. (1983) *Wizards of Armageddon*. Stanford: Stanford University Press.

Kapstein, E.B. (2011) Military metrics: How do we know when we're winning (or losing) a war? *Small Wars Journal*, Accessed at: <http://smallwarsjournal.com>.

Kerr, S. (1975) On the folly of rewarding A, while hoping for B. *Academy of Management Review* **18** (4), 769-783.

Kirkpatrick, C.E. (1992) *An Unknown Future and a Doubtful Present: Writing the Victory Plan of 1941*. Washington DC: US Army Center of Military History.

Koistinen, P.A.C. (2004) *Arsenal of World War II: The Political Economy of American Warfare, 1940-1945*. Lawrence, KS: University of Kansas Press.

Krepinevich, A.F. (1986) *The Army and Vietnam*. Baltimore: Johns Hopkins University Press.

Laffont, J., and Martimort, D. (2002) *The Theory of Incentives: The Principal-Agent Model*. Princeton: Princeton University Press.

Malcomson, J. M. (2009) Principal and expert agent. *The B.E. Journal of Theoretical Economics* **9** (1), Article 17 (Contributions).

McNamara, R.S. (1995) *In Retrospect: The Tragedy and Lessons of Vietnam*. New York: Times Books.

Millett, A.R., Murray, W. and Watman, K.H. (1986) The effectiveness of military organizations. *International Security* **11**(1): 37-71.

Mitnick, B.M. (1975) The theory of agency: The policing paradox and regulatory behavior. *Public Choice* **24** (1), 27-42.

Moskos, C.C. (1975) The American combat soldier in Vietnam. *Journal of Social Issues* **31** (4), 25-37.

Mueller, J.E. (1980) The search for the 'breaking point' in Vietnam: The statistics of a deadly quarrel. *International Studies Quarterly* **24** (2), 497-519.

Newland, S.J. and Chun, C.K.S. (2011) *The European Campaign: Its Origin and Conduct*. Carlisle Barracks, PA: Strategic Studies Institute.

North, D.C., Wallis, J.J. and Weingast, B.R. (2009) *Violence and Social Orders: A Conceptual Framework for Interpreting Recorded Human History*. New York: Cambridge University Press.

Overy, R. (1995) *Why the Allies Won*. New York: Norton.

Rosenau, W. (2005) *US Internal Security Assistance to South Vietnam*. London: Routledge.

Travis, D.S. (1990) US progress reports for the Vietnam War, 1967-68: A study of the Hamlet Evaluation System and the enemy order of battle. Unpublished thesis. Louisville, KY: University of Louisville.

Van Creveld, M. (1985) *Command in War*. Cambridge: Harvard University Press.

Weigley, R.F. (1973) *The American Way of War: A History of United States Military Strategy and Policy*. Bloomington, IN: Indiana University Press.

----- (1986) The Interwar Army, 1919-1941. In *Against All Enemies: Interpretations of American Military History from Colonial Times to the Present* edited by K.J. Hagan and W.R. Roberts. New York: Greenwood Press, 257-278.

Wirtz, J.J. (1991) Intelligence to please? The order of battle controversy during the Vietnam War. *Political Science Quarterly* **106** (2), 239-263.