



University of Warwick institutional repository: <http://go.warwick.ac.uk/wrap>

This paper is made available online in accordance with publisher policies. Please scroll down to view the document itself. Please refer to the repository record for this item and our policy information available from the repository home page for further information.

To see the final version of this paper please visit the publisher's website. Access to the published version may require a subscription.

Author(s): JQ Smith and L Dodd

Article Title: Regulating Autonomous Agents facing Conflicting Objectives: A Command and Control Example

Year of publication: 2011

Link to published article:

<http://www2.warwick.ac.uk/fac/sci/statistics/crism/research/2011/paper11-12>

Publisher statement: None

Regulating Autonomous Agents facing Conflicting Objectives: A Command and Control Example

Jim Q. Smith and Lorraine Dodd
Warwick University and Cranfield University

Abstract

UK military commanders have a degree of devolved decision authority delegated from command and control (C2) regulators, and they are trained and expected to act rationally and accountably. Therefore from a Bayesian perspective they should be subjective expected utility maximizers. In fact they largely appear to be so. However when current tactical objectives conflict with broader campaign objective there is a strong risk that fielded commanders will lose rationality and coherence. By systematically analysing the geometry of their expected utilities, arising from a utility function with two attributes, we demonstrate in this paper that even when a remote C2 regulator can predict only the likely broad shape of her agents' marginal utility functions it is still often possible for her to identify robustly those settings where the commander is at risk of making inappropriate decisions.

1 Introduction

To encourage its personnel to act as flexibly as possible, a command and control (C2) regulator (usually at a senior level of UK military command), devolves differing degrees of decision authority to field commanders. Within a mission statement, such a commander is given detailed information about how his decisive acts might provoke praise or inflame public opinion, promote or undermine trust and hence expedite or frustrate campaign objectives in the medium and long term. We henceforth say these factors frame the *strategic* objective. Such factors tend to encourage restraint, appeasement, the avoidance of conflict and embrace the need for public security. This objective will be complemented by the commander's particular mission task, called his *tactical* objective,

which directs him to achieve certain immediate military aims against an adversary.

The commander is trained and expected to act rationally. In particular, when given decision autonomy, he is aware that he will need to be able to justify his chosen acts in later after action reviews. In this paper a "rational" agent will be interpreted as someone who is a subjective expected utility maximizer.

Surprisingly, in contrast with some other domains, this type of Bayesian rationality assumption appears to be broadly consistent with how competent commanders usually act. Simulation experiments have identified only one scenario where such a commander's rationality is liable to be significantly undermined. This is when not only no act simultaneously appears to score well on both the set tactical and strategic objectives and but also two other relatively well scoring alternatives exist: one approximating being rational and ignoring the strategic objective whilst the other ignores the tactical objective. There is then a risk that a commander will exhibit various irrational behaviors such as hypervigilance [6] or decision suppression []. Even if the commander manages to remain rational then such scenarios will tend to induce discontinuity across adjacent autonomous commanders. For example, being faced with similar objectives and information, one rational commander may choose to commit to a full offensive whilst the other with similar training retreats - with potentially disastrous overall consequences. The C2 regulator therefore needs to address these risks in the way she communicates to her commanders.

One natural risk model for the C2 regulator to use would be to assume that her expected utility maximizing commander as having a utility with two attributes.- broadly measuring success in strategic and tactical objectives respectively. Her problem would then translate into one where she must minimize situations where her commander's expected utility exhibited two isolated local maxima, one favoring prioritizing the tactical objective and the other the strategic objective. However one apparent challenge in pursuing this approach is that the regulator is *remote*. Her information is therefore likely to inform her only about the broad characteristics of the geometry of the commander's utility function and posterior density. In this paper we demonstrate that the regulator can nevertheless still address these risks using only the coarse information plausibly available to her.

The development below may look speculative. However it is firmly based on evidence from on-going studies of decision-taking under simulated conditions of internal command contention and situational uncertainty, applied to the domain of military C2 on serving and experienced

military commanders see e.g. [1],[3]. In the next section we present a formal framework within which a regulator can address the risk of losing smooth C2 agility. We then present a new illustrative example of the type of scenario that can give rise to this loss of agility and how our formal approach can model this risk. After reviewing some pertinent general results from singularity theory [4] we then proceed to demonstrate how, under certain mild regularity conditions, the remote regulator can make broad risk assessments about the effectiveness of the decision making of her agent more generally.

2 Some Technical Apparatus

A rational commander decides an action $\mathbf{d}^* \in \mathbf{D}$ maximizing the expectation of his utility function U . Initially assume that a commander's utility function $U(\mathbf{d}, \mathbf{x})$ has two *value independent attributes* $\mathbf{x} = (x_1, x_2) \in \mathcal{X}_1 \times \mathcal{X}_2$ [5], [11] where x_1 measures the tactical success and x_2 the strategic success of the mission. Then we can write

$$U(\mathbf{d}, \mathbf{x}) = k_1 U_1(\mathbf{d}, x_1) + k_2 U_2(\mathbf{d}, x_2)$$

$\mathbf{d} \in \mathbf{D}$, $\mathbf{x} = (x_1, x_2) \in \mathcal{X}_1 \times \mathcal{X}_2$, where the i^{th} marginal utility $U_i(\mathbf{d}, x_i)$ is a function of its argument only and the *criteria weights* $k_i(\boldsymbol{\lambda}_1)$ satisfy $k_i(\boldsymbol{\lambda}_1) \geq 0$, $i = 1, 2$, $k_1(\boldsymbol{\lambda}_1) + k_2(\boldsymbol{\lambda}_1) = 1$, $i = 1, 2$. The rational commander then chooses a *Bayes decision* $\mathbf{d}^* \in \mathbf{D}$ to maximize his expected utility

$$\bar{U}(\mathbf{d}) = k_1 \bar{U}_1(\mathbf{d}) + k_2 \bar{U}_2(\mathbf{d}) \quad (1)$$

where for $i = 1, 2$

$$\bar{U}_i(\mathbf{d}) = \int U_i(\mathbf{d}, x_i) p_i(x_i) dx_i \quad (2)$$

and where $p_i(x_i)$ is his posterior marginal density of the attribute x_i , $i = 1, 2$.

Next express any course of action $\mathbf{d} = (d, \mathbf{d}_1, \mathbf{d}_2) \in \mathbf{D} = D \times \mathbf{D}_1 \times \mathbf{D}_2$ - for \mathbf{D} a possibly complicated decision - where D is a subset of the real line. The decision $d \in D$ represents the *intensity* of engagement of an associated action $\mathbf{d} \in \mathbf{D}$. Let the embellishment $(\mathbf{d}_1^*(d), \mathbf{d}_2^*(d))$ denote respectively the commander's decision to use an intensity d and then chooses a best action $\mathbf{d}_1^*(d)$ ($\mathbf{d}_2^*(d)$) to address respectively the tactical (strategic) objectives consistent with this chosen intensity of engagement. In many scenarios [4] it is possible to argue that, for a fixed embellishment $(\mathbf{d}_1^*(d), \mathbf{d}_2^*(d))$, a commander's expected tactical achievement $\bar{U}_1(\mathbf{d}_1^*(d), \mathbf{d}_2^*(d))$ is non-decreasing in his choice of intensity d whilst

his strategic achievement $\bar{U}_2(\mathbf{d}_1^*(d), \mathbf{d}_2^*(d))$ is non-increasing in d . Thus the higher his intensity of engagement the better the chances of tactical success but the worse the chances of strategic success. For simplicity assume that $\mathbf{d}_1^*(d)$ and $\mathbf{d}_2^*(d)$ do not to constrain one another given d so that the marginal utility $\bar{U}_i(\mathbf{d})$ is a function only of (d, \mathbf{d}_i) , $(d, \mathbf{d}_i) \in D \times \mathbf{D}_i, i = 1, 2$, where $D \subseteq \mathbb{R}$. Under these mild assumptions it is easily checked (see [4]) that a commander's Bayes intensity d^* will then maximize

$$V(d) = rP_1(d) - P_2(d) \quad (3)$$

where $P_1(d)$ ($P_2(d)$) are each distribution functions and respectively increasing (decreasing) linear transformations of $\bar{U}_i(d, \mathbf{d}_i^*(d))$ $i = 1, 2$.

Let $P_i(d)$ have support $[a_i, b_i]$, where, with an abuse of notation, we allow the lower bounds to take the value $-\infty$ and the upper bound ∞ . $i = 1, 2$. By this definition the commander believes that an intensity d below a_1 will surely achieve no tactical success but above b_1 he will surely be completely successful tactically. Similarly a_2 is the highest intensity he could use without damaging the strategic objective and the bound b_2 is the lowest value at which the strategic objective would be most severely compromised. Two particular decisions will be important reference points: an action with intensity $d = b_1$ we call *pure combat* and an action $d = a_2$ we call *pure circumspection*.

Let $u_i^- = \inf_{d \in D} \bar{U}_i(d, \mathbf{d}_i^*(d))$ and $u_i^+ = \sup_{d \in D} \bar{U}_i(d, \mathbf{d}_i^*(d))$ denote, respectively, the worst and best possible outcomes - for each of the objectives. We shall call the parameter

$$\rho \triangleq \log r = \{\log k_1 - \log k_2\} + \{\log (u_1^+ - u_1^-) - \log (u_2^+ - u_2^-)\}$$

the *daring*. Without loss we can assume the parameters of the problem are $(\rho, \boldsymbol{\lambda}) \in \mathbb{R} \times \Lambda$, where $P_1(d|\boldsymbol{\lambda}), P_2(d|\boldsymbol{\lambda}), \lambda \in \Lambda, i = 1, 2$ are chosen to be functionally independent of ρ . The first term in the above expression is increasing in the relative weight the commander places on tactical verses strategic objectives whilst the second measures the extent to which he expects he would attain his tactical objective were he to focus only on this less the extent he expects he could attain his strategic objective were he to address on that. Note that, because she is remote from the field of engagement a regulator may have difficulty accurately predicting her commander's chosen value of ρ in any given scenario.

The value of a commander's daring ρ impacts significantly on his decision making. As $\rho \rightarrow -\infty$ his expected utility will tend to his expected marginal utility on x_2 and so pure circumspection a_2 tends to optimality. On the other hand as $\rho \rightarrow \infty$ pure circumspection tends to optimality [4]. Furthermore if $d' < d$ and d' is not preferred to d when

$\rho = \rho_-$ then d' is not preferred to d when $\rho = \rho_+$, for any $\rho_+ \geq \rho_-$ and if $d' > d''$ and d' is not preferred to d'' when $\rho = \rho_-$ then d' is not preferred to d'' when $\rho = \rho_+$, for any $\rho_+ \leq \rho_-$. In this sense a rational commander's predisposition to act with intensity d is non-decreasing in his daring ρ *whatever* his circumstances. In particular, if his Bayes decision $d^*(\rho)$ is unique for all values of ρ for a fixed value of λ then it is non-decreasing continuously in ρ for that fixed λ (see [4]).

Rationality and contiguity challenges arise when two similar values of ρ can provoke the commander to choose two very different intensities of action. For then a commander's Bayes decision can suddenly can jump to a new much higher value even if ρ increases infinitesimally. So a commander facing a slight change in circumstances or confidence might suddenly regret not committing to a different decision. Furthermore in these circumstances two contiguous but autonomous commanders with similar expected marginal utilities could choose radically different intensities. So we study here when and how these risks might arise.

In fact the fairly weak assumptions above immediately imply three characteristics of the behavior of a rational commander:

- In [4] we prove that if $P_2(d)$ stochastically dominates $P_1(d)$ for feasible choices of intensity d .i.e. that

$$P_1(d) \leq P_2(d) \tag{4}$$

- for example if $b_2 \leq a_1$ - then the intensity chosen by any rational commander will be either pure combat or pure circumspection depending on his daring. Therefore the regulator can only address the discontiguity or contradiction induced by this *unresolvable* situation by trying to ensure that any commander's ρ are all of the same sign and well away from zero. So this scenario is always risky.

- When $b_1 \leq a_2$ then any “OK” courses of action (for example ones obtained from the Command Estimate Process) [7] with an associated intensity $d^* = [b_1, a_2]$ will be optimal. By choosing d^* the commander believes he will simultaneously achieve the maximum possible reward for both objectives. A *resolvable* situation has low risk: the autonomous commanders can be predicted largely to act appropriately. They will have been trained to search and find such OK actions.
- When $b_1 > a_2$ there is always a Bayes decision for the intensity which is pure combat b_1 or circumspection a_2 or some compromise intensity d , $\max\{a_1, a_2\} \leq d \leq \min\{b_1, b_2\}$: so henceforth assume the commander will choose his optimal intensity from this set.

Then when $a_2 < a_1$ or $b_2 < b_1$ the Bayes intensity will always be a discontinuous function of ρ see [4], hence always potentially risky.

- The finally case is where $a_1 \leq a_2 < b_1 \leq b_2$, when from the above, $d^* \in [a_2, b_1]$. Here the regulator has some control, so we study this case further below. Depending on how she communicates the mission her commander might choose an extreme decision close to pure circumspection or combat or some other intermediate intensity of engagement. We argue below that whenever possible a C2 regulator should try to construct scenarios where the commander will be predisposed to choose such an intermediate intensity because this will then expose her to less risk.

3 An Example of a Military Conflict Decision

In the following example we set $a_1 = a_2 = -1$ and $b_1 = b_2 = 1$, so that under the assumptions above the rational commander will choose an intensity $-1 \leq d \leq 1$, and illustrate the critical last scenario of the previous section.

Example 1 *A battle group is set the tactical objective of securing two districts a, b of a city. Its commander believes that each district will take a minimum time to clear plus an exponentially distributed delay x_a and x_b with rate parameter β_1 . He believes he will have failed tactically unless this task is fully completed. However before the securing a and b his strategic task is to first evacuate vulnerable civilians from two other districts c, d . There are potential delays x_c and x_d with rate parameter β_2 to add to the minimum time to complete this strategic objective and he believes he will have completely failed strategically unless he is able to evacuate both areas successfully. The whole mission must not be delayed by more than 2 units of time. The commander must commit now to the time $d + 1$ he allows for delays in the tactical objective (so implicitly budgets for a $1 - d$ delay in completing the evacuation. He believes that $\beta_1 \amalg \beta_2$ and that β_i has a gamma density*

$$\pi(\beta_i) = \frac{\lambda_i^{\alpha_i}}{\Gamma(\alpha_i)} \beta_i^{\alpha_i-1} \exp(-\lambda_i \beta_i), \quad > 0$$

where $\alpha_i, \lambda_i > 0$, so the delay he expects for each operation is $\mathbb{E}(\beta_i) = \alpha_i \lambda_i^{-1}$, $i = 1, 2$.

The sum of two independent exponential variables with the same rate β has a Gamma $G(2, \beta)$ distribution. Therefore his predictive density

for the delay experienced in each of the tasks has a density $p_i(t)$, $t > 0$, with a unique mode $(\alpha_i + 1)^{-1}\lambda_i$, given by

$$p_i(t) = \int_{\beta>0} \beta^2 t \exp(-\beta t) \frac{\lambda_i^{\alpha_i}}{\Gamma(\alpha_i)} \beta^{\alpha_i-1} \exp(-\lambda_i \beta) d\beta = \frac{\Gamma(\alpha_i + 2) t \lambda_i^{\alpha_i}}{\Gamma(\alpha_i) [t + \lambda_i]^{\alpha_i+2}}$$

$i = 1, 2$. Since his utility function is zero-one on each attribute, his expected utility associated with using intensity y , $-1 \leq y \leq 1$, is proportional to

$$rP_1(1 + y) + P_2(1 - y)$$

where $P_i(t)$ is the distribution function associated with $p(t)$, $i = 1, 2$. Any Bayes intensity must therefore satisfy

$$rp_1(1 + y) = p_2(1 - y)$$

For illustration first suppose $\alpha_1 = \alpha_2 = 1$. and $\sigma \triangleq r\lambda_1\lambda_2^{-1} = 1$. Then this equation rearranges to the cubic

$$ay(1 - y^2) + b(1 - y^2) + cy + e = 0$$

where $a \triangleq 4 + 3(\lambda_1 + \lambda_2)$, $b \triangleq 3[\lambda_1(1 + \lambda_1) - \lambda_2(1 + \lambda_2)]$, $c \triangleq [\lambda_2^3 + \lambda_1^3]$ and $e = \lambda_1^3 - \lambda_2^3$.

By letting $f \triangleq \frac{b}{3a}$, $g \triangleq (1 - a^{-1}c)$ and $z = y + f$ the equation above rearranges into

$$z^3 + \{9f^2 + g\}z - f(g - 2f^2 - e) = 0$$

The local maxima of the commander's expected utility function can therefore be described by the well studied canonical cusp catastrophe [12] [9] where the splitting factor of this cusp catastrophe is $-(9f^2 + g)$. So if

$$9f^2 + g \geq 0 \iff b^2 \geq a(a - c)$$

the commander's expected utility function can have only one local maxima: a "compromise" Bayes intensity. When parameters summarizing the commander's information and his values lie in this region his choice of optimal decision will therefore be a smooth function of those parameters: similarly trained and tasked adjacent commanders will tend to act similarly, and small changes in the commander's information will not lead to large changes in what he perceives to be his optimal decision.

On the other hand if

$$9f^2 + g < 0 \iff \frac{b^2}{a^2} < 1 - \frac{c}{a} \iff b^2 < a(a - c)$$

then, for values of $f(g - 2f^2 - e)$ close to zero his expected utility function will have two local maxima - one nearer pure circumspection and the other nearer pure combat with a local minimum in between. Small changes in parameters may then cause different contiguous commanders to choose different decisions or a single commander to regret having committed to his chosen act.

So for example in the completely symmetric scenario when $\lambda_1 = \lambda_2 = \lambda$ there is always a stationary point at 0. When λ is very large (so that the expected delays are very small) 0 is the unique stationary point. However the "compromise" $d = 0$ is actually a *minimum* of the expected utility - i.e. locally the poorest - whenever

$$2 + 3\lambda - \lambda^3 > 0$$

It is easy to check that this inequality is satisfied only if $\lambda^{-1} > \frac{1}{2}$ i.e. when the sum over the four expected delays is greater than the total time allowed for the mission 2. The larger this expected total delay is the more extreme the two equally preferable alternative decisions are. Clearly then the regulator should try to ensure that her commander is not faced with a scenario where $9f^2 + g < 0$, when in the case $\lambda_1 = \lambda_2 = \lambda$ translates into the commander has been given enough time to expect to complete both parts of his mission.

4 The General Case

The example above would be unremarkable if the geometry of the expected utility were heavily dependent on the exact form of distributional assumptions made by the commander. However it can be shown that *qualitatively* the geometry determining whether or not a given scenario exhibits these risks is surprisingly robust to changes in the algebraic forms of the quantitative forms of the commander's beliefs and values. So the sorts of broad criteria that a regulator should adopt, such as allowing the commander enough time to expect he can complete his mission successfully, can be determined in a much more general framework. In this general framework the conditions needed to avoid bipolar decision making are expressed in terms of points of inflexion of certain functions. Just as for the example above any given scenario the position of these points can in turn usually be expressed in a qualitatively meaningful way and subject to the influence of the regulator.

Thus using [7], [8], [10], we next investigate the geometrical conditions determining when such risks might exist. Henceforth assume that the regulator believes that her commanders all have distributions P_i that are unimodal and twice differentiable in the open interval (a_i, b_i) , $i = 1, 2$ and constant nowhere in this interval and fix the value of the parameters

λ . Any local maximum of $V(d)$ will then either lie on the boundary of the feasible space or satisfy

$$v(d) \triangleq f_2(d) - f_1(d) = \rho \quad (5)$$

where $f_i(d) = \log p_i(d)$, $i = 1, 2$. Provided the derivative $Dv(d) \geq 0$ such a stationary point must be a local maximum of V .

Let ξ_i denote the maximum (or mode) of $p_i(d)$ $i = 1, 2$. Because ξ_1 is a point of highest incremental gain in mission we call this point the *mission point* and the intensity ξ_2 where the threat to campaign objectives worsens fastest the *campaign point*. Note that it is not unreasonable to assume that two similarly trained and missioned commanders will entertain similar campaign and mission points when facing similar scenarios.

From this definition if the modes satisfy $\xi_1 \leq \xi_2$, then, for any $d \in [\xi_1, \xi_2]$, $v(d)$ is strictly decreasing. It follows that there is at most one solution d^* to (5) for any value of ρ and $Dv(d) \geq 0$. Therefore this stationary value $d^* \in (a, b)$ is a local maximum of V . When $\xi_1 \leq \xi_2$ for each value of $\lambda \in \Lambda$: so that the mode of the intensity of tactical success is less than the modal intensity for strategic success, there is a *unique* interior maximum in this interval moving as a continuous function of λ . Consequently a regulator will find such a scenario almost as desirable as resolvable ones where the actions of the devolved decision-takers work well. Although their actions will depend on ρ , two commanders with similar but different daring ρ will act similarly and in particular it is rational for them both to compromise. So if contiguous commanders are matched by their training and emotional history then they will make similar and hence broadly consistent choices. In our running example this scenario would correspond to one where the modes of the sum of the two delay distributions was less than the time allowed for the this delay.

On the other hand we prove in [4] that when $\xi_2 < \xi_1$ there is a value of ρ and an associated decision which is a stationary point of V for which the derivative of $v(d)$ is negative so that such a stationary point will be a local *minimum*. Then, just as in our example, the set of optimal decisions bifurcates into two disjoint sets: one lying in the interval of a "lower intensity" consistent with strategic objectives, and the other in an interval of intensity favouring achieving mission objectives. At value of $\rho = \rho^\circ$ an option in each of these sets will be optimal giving rise to bipolarity in the decision space.

Example 2 *In our running example but when all prior parameters are arbitrary the condition on the two modes ξ_1, ξ_2 above implies bifurcation can occur for some value of ρ iff*

$$(\alpha_1 + 1)^{-1} \lambda_1 + (\alpha_2 + 1)^{-1} \lambda_2 < 2$$

The original interpretation that risks are avoided if the delays are perceived as not too large is therefore retained. When $\alpha_1 = \alpha_2 = \alpha$ and $\lambda_1 = \lambda_2 = \lambda$ this condition says that the prior expectation μ for dealing with any area satisfies $\mu \leq (1 + \alpha)^{-1}$ where $\alpha^{-1/2}$ is the coefficient of variation of the prior. So as $\alpha \rightarrow \infty$ there is less and less point aborting the evacuation because the commander believes he is unlikely to encounter any less delays in the tactical elements of his mission.

From the geometrical arguments given above we can argue that the phenomena we illustrated with our last example are generic and depend only on certain broad features such as that the predictive distributions associated with the two attributes being unimodal and smooth. It can even be shown that even the unimodality conditions can be relaxed. It follows that the regulator can be fairly confident of the types of scenarios where his agents may not act well. Of course she will not be able to predict the precise values of the critical parameters like the sign of $\xi_1 - \xi_2$. But in many situation she can broadly estimate this using the information she has at hand or from general experiential knowledge.

5 Robustness and Links with Game Theory

The running example could be extended into a problem in dynamic stochastic control where the commander has a choice of whether or not to abort the evacuation at *any* given time and simply focus on the tactical imperatives. Although the associated technicalities are beyond the scope of this paper, an optimal policy is straightforward to calculate. The qualitative behavior of the commander is analogous to the one above: the only difference is that he tends to give up earlier on the evacuation after he learns that the delays to this operation turn out to be longer than predicted.

However, much more critical to the qualitative form of our deductions is the commander's prior beliefs about the dependence of the parameters of his distribution and the actual form of utility independence he holds. Suppose he believes that the delays he incurred were not due to the nature of the task but the competence of his unit to accomplish *any* task speedily. Then in the dynamic control setting above, instead of believing $\beta_1 \perp \beta_2$ he might believe that $\beta_1 = \beta_2$. Were both his criterion weights $\frac{1}{2}$ then it is easily checked that pure circumspection is always optimal for him: i.e. he should continue evacuation until the two area are clear and only then attempt the tactical objective. Of course when a regulator is familiar with the likely qualitative form of the *reasoning* used by her commander, here whether he assumes independence or identity, then his behavior will still be qualitatively predictable to her.

Another qualitatively different setting, potentially giving rise to different behavior is when the commander does not have value independent attributes. For example suppose his mission of securing areas c and d will have failed unless he has first evacuated areas a and b . In this case his utility function will be of the form $U = k_1U_1U_2 + k_2U_2$. Then obviously pure circumspection - "continue until evacuation is finished and only then secure the other areas" is again always optimal. A regulator will thus incline contiguity by communicating the mission in this way. Again the difference in qualitative form of the utility should be relatively predictable to the regulator.

Finally in these continuous models of decision process by autonomous agents a regulator can explore efficacious ways to try to influence the actions of her remote agent by partitioning the parameter space into *finite discrete* subsets within which certain qualitative features of the agent's behavior are predictable. This finiteness opens the possibility of drawing on game theory to further our understanding of such scenarios. Note that the contiguity of the opponent's rationality may be undermined by deceiving the adversary into facing its commanders with bipolarity. Thus in our running example if mission statements are consistent with only short delays being experienced by a commander then the regulator will be confident she is in a situation which is OK and controllable on the ground. Therefore if an adversary can deceive her into being over optimistic about these delays then her operations are liable to be more easily disrupted. Similarly, if a confident commander can be frustrated early in his actions so that he becomes more uncertain - for example by the adversary throwing all its effort into frustrating his evacuation of area a in the example above - then that commander may again be drawn into bipolar decision making, potentially undermining his rationality and therefore his later effectiveness. As far as we know the study of these sorts of developments, combining experimental psychology, game theory, singularity theory and dynamic stochastic control are in their infancy. In a later paper we will communicate some of our early results about these synergies as they apply to this domain.

References

- [1] Dodd, L. Moffat, J Smith, J.Q. and Mathieson. G.(2003) "From simple prescriptive to complex descriptive models: an example from a recent command decision experiment" Proceedings of the 8th International Command and Control Research and Technology Symposium June Washington
- [2] Dodd, L.(1997), "Command decision studies for future conflict" DERA Unpublished Report.

- [3] Dodd, L. Moffat, J. and Smith, J.Q.(2006) "Discontinuities in decision-making when objective conflict: a military command decision case study" *J.Oper. Res. Soc.*,57, .643 - 654
- [4] Dodd, L. and Smith, J.Q. (2010) "Devolving Command Decisions in Complex Operations" *CRISM Res. Rep.* 10 -17
- [5] French, S. and Rios Insua, D.(2000) "Statistical Decision Theory" Arnold
- [6] Janis, J.L. and Mann, L.(1977) "Decision Making: A Psychological Analysis of Conflict, Choice and Commitment" Free Press. N.Y.
- [7] Moffat, J. (2002) "Command and Control in the Information Age" The Stationary Office, London
- [8] Moffat,J. and Witty, S. (2002) "Bayesian Decision Making and military control" *J. Oper. Res. Soc.* , 53, 709 - 718
- [9] Poston, T. and Stewart, I. (1978) "Catastrophe Theory and its applications" Pitman
- [10] Smith, J.Q., Harrison, P.J., and Zeeman, E.C.(1981) "The analysis of some discontinuous decision Processes" *E. J. Oper.Res.* Vol. 7, 30-43.
- [11] Smith, J.Q.(2010) "Bayesian Decision Analysis: Principles and Practice" Cambridge University Press
- [12] Zeeman E.C.(1977) "Catastrophe Theory: Selected Papers" Addison Wesley