

This article was downloaded by: [Markman, Arthur]

On: 28 March 2011

Access details: Access Details: [subscription number 935704877]

Publisher Taylor & Francis

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Biological Rhythm Research

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t713734219>

The effects of 24-hour sleep deprivation on the exploration-exploitation trade-off

Brian D. Glass^a; W. Todd Maddox^{ab}; Christopher Bowen^c; Zachary R. Savarie^c; Michael D. Matthews^c; Arthur B. Markman^a; David M. Schnyer^{ab}

^a Department of Psychology, University of Texas at Austin, Austin, TX, USA ^b Institute for Neuroscience, University of Texas at Austin, Austin, USA ^c Department of Behavioral Sciences, United States Military Academy, West Point, NY, USA

First published on: 07 July 2010

To cite this Article Glass, Brian D. , Maddox, W. Todd , Bowen, Christopher , Savarie, Zachary R. , Matthews, Michael D. , Markman, Arthur B. and Schnyer, David M.(2011) 'The effects of 24-hour sleep deprivation on the exploration-exploitation trade-off', Biological Rhythm Research, 42: 2, 99 – 110, First published on: 07 July 2010 (iFirst)

To link to this Article: DOI: 10.1080/09291011003726532

URL: <http://dx.doi.org/10.1080/09291011003726532>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

The effects of 24-hour sleep deprivation on the exploration–exploitation trade-off

Brian D. Glass^{a,*}, W. Todd Maddox^{a,b}, Christopher Bowen^c,
Zachary R. Savarie^c, Michael D. Matthews^c, Arthur B. Markman^a and
David M. Schnyer^{a,b}

^aDepartment of Psychology, University of Texas at Austin, Austin, TX, USA; ^bInstitute for Neuroscience, University of Texas at Austin, Austin, USA; ^cDepartment of Behavioral Sciences, United States Military Academy, West Point, NY, USA

(Received 22 January 2010; final version received 18 February 2010)

Sleep deprivation has a complex set of neurological effects that go beyond a mere slowing of mental processes. While cognitive and perceptual impairments in sleep deprived individuals are widespread, some abilities remain intact. In an effort to characterize these effects, some have suggested an impairment of complex decision-making ability despite intact ability to follow simple rules. To examine this trade-off, 24-hour total sleep deprived individuals performed two versions of a resource acquisition foraging task, one in which exploration is optimal (to succeed, abandon low value, high saliency options) and another in which exploitation is optimal (to succeed, refrain from switching between options). Sleep deprived subjects exhibited decreased performance on the exploitation task compared to non-sleep deprived controls, yet both groups exhibited increased performance on the exploratory task. These results speak to previous neuropsychological work on cognitive control.

Keywords: sleep deprivation; exploration-exploitation; reasoning; decision making; resource acquisition; foraging

Introduction

Sleep deprivation (SD) is a way of life for many people who perform critical societal roles such as those in the military, and those in medical professions (Belenky et al. 1994). Therefore, it is important to have a thorough understanding of the role of sleep deprivation on reasoning and decision making. The classical view of the effects of sleep deprivation is one of an overall slowing of cognitive abilities. This view has been called into question, because not all cognitive abilities are equally affected by sleep deprivation (Harrison and Horne 2000). In a battery of tasks, Williamson et al. (2000) found that perceptual, attentional and memory tasks were impaired by sleep deprivation, but visual search and logical reasoning tasks were not. Maddox et al. (2009) found deficits in perceptual classification ability in an information-integration task for some sleep-deprived

*Corresponding author. Email: glass@mail.utexas.edu

individuals, but not all. Herscovitch et al. (1980) found that sleep deprivation led to increased perseveration on the Wisconsin Card Sorting Task (WCST), indicating a decreased ability to shift set and try new strategies. Neuropsychological evidence points towards decreased activation of the prefrontal cortex (PFC) after sleep deprivation, leading to decreased inhibitory control (Chuah et al. 2006; Tomasi et al. 2008; Maddox et al. 2009).

Other studies revealed risk-seeking behaviour in sleep deprivation, especially in known-risk tasks (Neri et al. 1992; Harrison and Horne 1998; McKenna et al. 2007). In a review of the sleep deprivation literature, Harrison and Horne (2000) endorsed the view that sleep deprivation has little effect on simple rule-following tasks, but it impedes decision making in complex integration tasks requiring flexibility, innovation or plan revision. However, other research shows lapsing and performance decrements in relatively simple tasks such as identification and vigilance tasks (Chee et al. 2008; Ratcliff and Van Dongen 2009). Other investigations have demonstrated increased inhibition and decreased cognitive control after sleep deprivation. For example, Drummond et al. (2006) and Chuah et al. (2006) demonstrated decreased inhibition in a Go-NoGo task. In a MRI study, Venkatraman et al. (2006) found elevated expectation of higher reward and diminished responses to losses in a gambling task. Interestingly, Drummond et al. (2001) and Strangman et al. (2005) conclude that the brain may be more plastic after sleep deprivation than previously thought by showing compensatory brain activation depending on task demands. Some of these results appear contradictory, with some showing evidence of decrements in simple tasks and others demonstrating intact performance in simple tasks despite decrements in complex tasks.

We examine another view of the influence of 24-hour total sleep deprivation on cognitive processes by focusing on the trade-off between exploration and exploitation. *Exploration* is the willingness to try new strategies and options, while *exploitation* is the reliance on past knowledge and known successful strategies. Adopting a new strategy requires a momentary or permanent abandonment of prior strategies, resulting in a trade-off (Daw et al. 2006; Worthy et al. 2007). While previous work has been done on parallel constructs such as risky behaviour and rule-based decision making, the exploration–exploitation trade-off offers a unique framework for study of the cognitive effects of sleep deprivation.

In this paper, we used a resource acquisition foraging task to examine the impact of sleep deprivation on the exploration–exploitation trade-off. In this task, participants compete for resources against computerized opponents in a real-time 2D simulation (see Roberts and Goldstone 2006). Participants are confronted with the exploration–exploitation trade-off in an attempt to maximize resources acquired. In some versions, exploration is optimal (i.e., maximizes performance), while in others exploitation is optimal.

Here we provide a brief introduction to the task and sleep deprivation procedure. The task environment is comprised of a large computerized grid, for which spaces can be occupied by the participant's token, the tokens of nine computerized opponents, or tokens that represent resources which the participant and opponents attempt to collect. The participant moves his or her token around the grid using a keyboard, while the opponents are driven by a simulation model described below. Resources appear steadily throughout the task, but differ in their placement on the grid depending on whether the exploitative or

exploratory task environment is in effect. In the exploitative task, resources appear in two patches, or clumps, on different sides of the grid (Figure 1a). The optimal strategy is to remain at one patch; switching patches wastes valuable time and will prevent the participant from reaching a bonus criterion (described below). In the exploratory task, there is a large central patch comprised of low value resources (result in fewer points). In small patches on the edges of the grid, there are higher value resources which must be found in order to reach the bonus criterion (Figure 1b). To ensure participants must explore for these peripheral resources, the board is hidden except for a small radius around the participant's token. This ensures that exploration of the space is required to succeed.

In the exploitative task, a key measure of performance is the number of patch switches (fewer is optimal). In the exploratory task, the measure is the distance from the centre of the grid (larger distances are more optimal since the higher valued patches are farther from the centre of the grid). In both tasks, resource score is a general performance measure. In the current study, all participants completed two morning sessions (comprised of three exploitative and three exploratory tasks) separated by 24 hours. The sleep-deprived group did not sleep between sessions, while the control group slept normally. Surface features are changed between runs, such as background colour and token colour and location of the patches, but the basic structure of the tasks remained unchanged. For example, in one run the high replenishment patch is in the upper right and the low replenishment patch is in the lower left, while in another run the locations are rotated and reversed.

Taken together, previous behavioural and neuropsychological work in sleep deprivation predicts impaired performance on the exploitative tasks, due to a decreased ability to continually apply a simple rule. It is unclear whether performance in the exploration tasks will be facilitated. In the following sections we




Figure 1. (A) Example of an exploitative environment simulation after 750 time steps. The lone participant agent is represented as a square with a white X, opponent agents are represented as squares, and resources are represented as circles. (B) Example of an exploratory environment simulation after 750 time steps. Only resources and other agents within 20 grid spaces are visible to the agents (represented by the circle around the lone participant agent). The resources in the central patch have a low point value, while the resources in the peripheral patches have a high point value.

introduce the model underlying these foraging tasks. Finally, we describe an experiment comparing the performance of sleep deprived and control subjects.

Methods

Sleep deprivation

For this experiment, both groups were non-sleep-deprived on day 1, and a day 2 sleep-deprivation manipulation was between subjects. All subjects were tested between 06:00 am and 11:00 am at the same time both days. Participants completed various tasks in an ongoing sleep study, with breaks throughout. Alcohol consumption was prohibited in the 24 hours prior to day 1, as was caffeine between 12:00 am and 6:00 am and during task participation. Participants were instructed to engage in normal sleep-wake cycles prior to day 1 testing, and engaged in peer monitoring during this period. After day 1 testing, a research assistant monitored the sleep-deprived group from 7:00 pm on day 1 until monitored testing at 6:00 am on day 2. Overnight, the sleep-deprived group ate dinner and engaged in activities such as walking, bowling, and video and board games. The control group was instructed to resume normal sleep. All testing took place in a university building.

Design

Participants engaged in the exploitation and exploration tasks on day 1 and 2. This resulted in a 2 (task) \times 2 (day) \times 2 (group) design. Each day, participants completed three runs of each task type, all with a separate bonus resource acquisition score criterion (described below). Resource acquisition score was the main performance measure for both tasks. Also, for the exploitation task, patch switches were measured. One patch switch was recorded for each movement from one patch to the other, defined as crossing the bound that bisects a line connecting the patch centres. Participants attempted to maximize score, and minimizing patch switches is the ideal strategy. For the exploratory task, the mean distance from the grid centre (i.e., the main central patch) was measured. Exploring the space and avoiding the central patch is the ideal strategy to maximize score.

Participants

Forty-one undergraduate cadets from the United States Military Academy at West Point participated in the study, with 23 (two female) non-sleep deprived controls and 18 (six female) sleep deprived. There was a slight difference in group age ($M = 18.6$, $SD = 0.8$ for controls; $M = 20.7$, $SD = 1.9$ for the sleep deprived group). Age as a covariate was never significant in the behavioural results. All participants had normal or corrected to normal vision. The Institutional Review Board of The University of Texas, Austin and West Point Military Academy approved the study and informed consent was obtained from all participants.

Materials

The task was run on a PC-compatible computer. Participants controlled their agent with the keyboard number pad. The participant agent and simulated agents could move horizontally, vertically and diagonally.

Procedure

On days 1 and 2, subjects performed three exploratory runs and three exploitative runs, interleaved, with each lasting 7.5 minutes.

In the exploitative task, participants gained 27 points for collecting a resource. Additionally, they gained one point when an opponent collected a resource. This small gain in points for the resources collected by competitors was instituted to control for effects of task reward structure on exploration that we have observed in previous studies (e.g., Worthy et al. 2007). The bonus criterion was set to 4037 (10% of all available resources points).

In the exploratory task, low-valued resources were worth 27 points and high-valued resources were worth 2700 points. Additionally, they gained points when opponents collected resources (one point for low-valued resources, 100 points for high-valued resources). The bonus was set to 17,950 (10% of all available resource points). Participants were told the scoring procedures and that both low- and high-valued resources were available. A point meter tracked points during the run. They were notified whether they achieved the bonus criterion after each run. No additional information was given between sessions, and instructions were reviewed on day 2.

For motivation to do well, participants had the opportunity to receive T-shirts for good performance. After day 2, for each subject, a random run was chosen from each day’s runs. If they reached the criterion score on that run then they received the T-shirt. This procedure was adapted from work on motivational influences on decision making. It was introduced in order to instantiate a promotion-gains state shown to encourage a polarizing motivational environment that minimizes the influence of inter-subject differences in motivational state (Worthy et al. 2007).

Surface features of the GEM environment were adjusted between runs. The colour of the participant agent, opponent agents, resources and background changed after each run. Also, different (but equivalent) reflections and rotations of the patch locations were used between runs.

Generalized Exploration Model (GEM)

GEM is an agent-based model of resource acquisition based on the EPICURE model of human foraging (Roberts and Goldstone 2006). A group of agents compete for resources by travelling around a 120 × 120 grid environment. A resource is represented by a static token that takes up one grid space and is removed only when an agent acquires it by landing on it. Each agent is driven by a separate instantiation of GEM. On each time step, each agent computes the expected value of the resource options available using Equations (1) through (5), where $\sigma = 4$, and (i, j) represent environment coordinates.

$$Value(i, j) = (P_1^* resourcedistance) - (P_1^* agentdensity) + (P_4 + goaldistance) \tag{1}$$

where,

$$resourcedistance(i, j) = \frac{e^{-distance^2/2\sigma^2}}{\sigma\sqrt{2\pi}} \tag{2}$$

$$agentdensity(i, j) = \sum_{otheragents} \frac{e^{-distance^2/2\sigma^2}}{\sigma\sqrt{2\pi}} \tag{3}$$

Downloaded By: [Markman, Arthur] At: 13:08 28 March 2011

for distance between (i, j) and nearby agent

$$goaldistance(i, j) = \frac{e^{-distance^2/2\sigma^2}}{\sigma\sqrt{2\pi}} \quad (4)$$

for distance between (i, j) and current goal location

$$distance(i, j) = \max(abs(x - i), abs(y - j)) \quad (5)$$

for all distance calculations.

In Equation (1), an expected value is computed for a given resource option as a linear combination of the distance to the resource, the density of agents around the resource, and the distance of the resource to the currently active goal resource location. Equations (2) and (4) describe the preference of shorter over longer distances via a Gaussian drop-off process. Equation (3) computes the density of nearby agents, also with a Gaussian process which emphasizes nearby agents. Equation (5) computes grid distance between locations. The same equations are used in both the exploitative and exploratory tasks. The expected values are re-evaluated at each time step.

After computing expected values, each agent selects a goal resource to travel towards using a soft-max decision process driven by an exploration parameter (K in Equation (6); see Daw et al. 2006; Worthy et al. 2007). The higher the exploration parameter, the less likely the agent will select the goal resource with the highest expected values, and thus act exploratory. The lower the exploration parameter, the more likely the agent will seek the highest expected value, thus acting exploitatively.

$$P(\text{moving to } i, j) = \frac{e^{\text{Value}(i,j)/K}}{\sum_x \sum_y e^{\text{Value}(x,y)/K}} \quad (6)$$

Typically, resources are allocated to the environment according to some scheme, such as ‘‘patches’’ of replenishing resources. A number of important experimental factors can be manipulated in this environment such as the locations of the patches, the replenishment rates of those patches, and the point values given to resources in different patches.

This environment is sufficiently complex and dynamic that simulations are required to determine the optimal behaviour of an agent in any given environment. For the present studies, we focus on environments with a single participant agent competing against nine opponent agents. To determine the strategy association with optimal behaviour in a given environment, the parameters of the opponent agents are held constant. In the simulations, we manipulated the exploration parameter values as well as the other parameter values of the single participant agent systematically to determine the optimal level of exploration in that environment. Using these simulations, we identified one environment in which the optimal behaviour for the participant is to exploit known resources and a second environment in which the optimal behaviour is to explore.

In the task for which exploitation is optimal, resources are dropped in two patches centred on locations 40 units apart (see Figure 1B). Resources are dropped in patch locations based on a random bivariate Gaussian distribution, to give the impression of a cluster. Resources appeared at a rate of 1 per 200 msec. The patches are replenished at different rates, so that 80% of the resources are dropped at the

high-rate patch and 20% at the *low-rate* patch. All resources in this task are worth a small number of points. Because resources are concentrated at the two patch locations, the ideal behaviour is to limit travel between patches. In simulations, agents with low exploration parameter values switch patches less often and acquire more resources. Abandoning a seemingly exhausted patch is suboptimal because patches continue to replenish and do not change locations. Also, competitors do not favour one patch over the other, and instead match their group allocation to the relative replenishment rates of the patches. This distribution further reinforces the benefit of not spending time patch switching.

In the task for which exploration is optimal, each agent has a limited sight radius of the total environment (as the highlighted circle in Figure 1B). They can detect resources and other agents within a radius of 20 grid spaces. This environment features low-valued resources worth a small number of points and high-valued resources worth many points. The low-valued resources are dropped steadily throughout the task at a patch located at the centre of the environment. High-valued resources appear sporadically in patches randomly located on the edges of the environment. The location of these patches changes randomly throughout the run. Therefore, to discover the high valued resources, foragers must abandon the plentiful yet low-valued central patch and explore the space. In simulations, agents with low exploration parameter values spend more time closer to the central pool and therefore acquire fewer resources. Thus, exploration is optimal in this task.

Behavioural experiments can be conducted by letting a human participant control one agent (with a directional keypad) while letting GEM control the other agents. The behavioural characteristics of the human participant can then be compared to those of the optimal forager. Pilot work with these environments suggested that participants are typically able to adapt their strategy to the task. So, participants in an exploitative environment learn to exploit a resource pool and participants in an exploratory environment learn to explore the environment.

In addition to using the number of resource points acquired as a gross measure of behaviour, it is possible to do more fine-grained analyses. For the exploitation task, we count the number of switches between patches. For the exploration task, we measure the distance from the centre of the grid. We expect subjects to have higher measures of exploratory performance when sleep deprived than when rested. This prediction is based on previous research which has demonstrated that sleep deprivation effects parallel constructs to exploration (e.g., decreased cognitive control and risky behaviour after sleep deprivation). However, in the end, the exploration–exploitation trade-off is not the same as those other constructs.

Results

For each task performed, we measured the number of resource points obtained. We also measured behavioural aspects that correlated with task performance in simulated agents.

Points obtained

An analysis compared day 1 performance between groups. Figures 2 and 3 represent the mean resource acquisition score on the exploitative and exploratory tasks on days 1 and 2. Importantly, there were no significant differences between the control




Figure 2. Means and standard errors of resource scores on the exploitative task for the control and sleep-deprived groups. The bonus criterion was set to 4037 points and is demarcated by the dashed line.

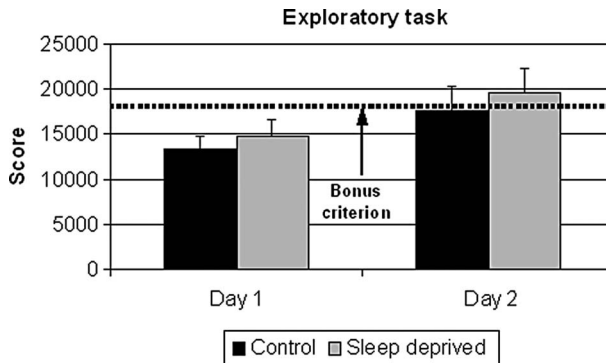


Figure 3. Means and standard errors of resource scores on the exploratory task for the control and sleep-deprived groups. The bonus criterion was set to 17,950 points and is demarcated by the dashed line.

and sleep-deprived groups on day 1 for the exploitative task, $t = 1.1$, $df = 37$, $p = 0.28$, nor the exploratory task, $t = -0.64$, $df = 37$, $p = 0.53$. This was expected since neither group was sleep deprived at this point. Since there was no overall difference on day 1, we conducted the analyses in two ways: averaged over the three runs on each day, and for the last run of day 1 and the first run on day 2.

We performed a repeated measures 2 (group) \times 2 (day) ANOVA on points obtained for the exploitation and exploration tasks separately. For the exploitation task, there was no main effect of day, $F(37) = 0.52$, $p = 0.48$, $\eta^2 = 0.014$. However, there was an interaction of group and day, $F(37) = 9.12$, $p = 0.005$, $\eta^2 = .0198$, and a main effect of group $F(37) = 9.90$, $p = 0.003$, $\eta^2 = 0.211$, such that the sleep-deprived group obtained fewer points on day 2 than the control group.

For the exploration task, there was a main effect of day, $F(37) = 11.41$, $p = 0.002$, $\eta^2 = 0.236$, such that both groups performed better on day 2. There was no main effect of type, $F(37) = 0.35$, $p = 0.557$, $\eta^2 = 0.009$, and no significant interaction of day and type $F(37) = 0.025$, $p = 0.876$, $\eta^2 = 0.001$.

On day 2, the resource acquisition scores for the sleep-deprived group were lower than the control group on the exploitative task, $t = 3.2$, $df = 22$, $p = 0.004$, and not significantly different than the control group on the exploratory task. Thus sleep deprivation led to a significant performance decline when exploitation was optimal, but not when exploration was optimal.

We repeated the above ANOVA using the third run from day 1 and the first run from day 2, instead of a combined score, and found a similar pattern. For the exploitative task, there was no main effect of day, $F(37) = 1.213$, $p = 0.278$, $\eta^2 = 0.032$, and no main effect of group, $F(37) = 1.50$, $p = 0.228$, $\eta^2 = 0.039$. However, day and group interacted significantly, $F(37) = 6.19$, $p = 0.017$, $\eta^2 = 0.143$, such that the control group obtained higher point scores. In the exploratory task, no effects were significant. In other words, on day 2 the sleep-deprived group performed worse than the control group in the exploitation task, but not in the exploratory task.

Foraging behaviour

In addition to the resource points obtained, movement behaviour sheds light on foraging strategies used by participants. In the exploitative task, computer simulations reveal that patch switches correlate negatively with resource acquisition. Switching patches is a time-consuming activity during which resources cannot be acquired. In the exploratory task, maintaining a greater distance from the environment centre is correlated positively with resource acquisition. The higher-value resources were located on the periphery, while a distracting low-value patch remained at the centre of the environment. Exploratory behaviour involved abandoning the central patch, thus distance from the centre of the central patch is an indicator of exploratory behaviour. For both tasks, there were no differences between the control and sleep-deprived groups on day 1.

In the exploitative task, we conducted a repeated measure 2 (group) \times 2 (day) ANOVA on patch switches. There was no main effect of day or significant interaction of day and group, but there was a nearly significant main effect of group $F(37) = 3.94$, $p = 0.055$, $\eta^2 = 0.096$ (Figure 4). Over all runs, the sleep-deprived group switches patches significantly more often than the control group on day 2,

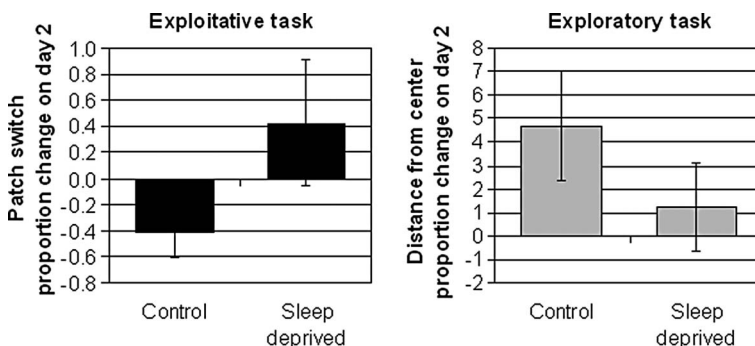


Figure 4. Means and standard errors of the number of patch switches in the Exploitative task and the distance from the centre of the environment in the exploratory task on day 1 to day 2 for the control and sleep-deprived groups.

$t(92) = -2.26, p = 0.026$. In the exploratory task, both groups travelled farther from the centre of the environment on day 2, with no significant differences between groups. Correlations verified the relationship between these foraging behaviour measures and resource score. Patch switches correlate negatively with exploitative task resource score, $R^2 = -0.413, p < 0.0001$. Distance from the grid centre correlated positively with exploratory task resource score, $R^2 = 0.273, p < 0.0001$. Scores between the tasks did not correlate significantly, nor did patch switches correlate with distance from centre.

Discussion

We found that sleep-deprived participants were more likely to explore an environment in a foraging task regardless of the optimal strategy in that environment. Sleep-deprived subjects were more exploratory on day 2 on both tasks, while control subjects were able to match their level of exploration to task demands. These results suggest that sleep deprivation leads to an increased tendency to abandon rules and adopt a more fluid and exploratory approach to the environment. This result paints a more complex picture than the hypothesis that sleep deprivation worsens overall performance in decision making tasks. This is not surprising given that complex decision making tasks likely require a number of components working together in concert. More aspects of this complexity will be revealed as the field continues an ever more fine-grained analysis of the decision making deficits associated with sleep deprivation.

Looking at performance more specifically, control subjects improved on both tasks on day 2. In contrast, sleep-deprived subjects were more exploratory in both tasks. Thus, on day 2 they exhibited worse exploitative task performance and improved exploratory task performance.

The results may provide insight into the effect of sleep deprivation on cognitive control and its role in inhibiting suboptimal strategies. Previous neuropsychological work observed impaired cognitive control associated with decreased PFC activity after sleep-deprivation (Chuah et al. 2006; Tomasi et al. 2008). Ishii et al. (2002) proposed that PFC, known for a directing role in decision-making, behavioural inhibition, and planning, also directs cognitive control of the exploration–exploitation trade-off in the evaluation of dynamic task environments. Our results suggest that impaired cognitive control or increased inhibition of the maintenance of the exploration–exploitation trade-off could have led to the sleep-deprived group's increased exploration. On day 2, they were less exploitative in the exploitative task and more exploratory in the exploration task, resulting in a significant interaction of group and task. Therefore, the results are consistent with previous findings of decreased inhibition after sleep-deprivation and further provide evidence that this inhibition may lead to a tendency toward exploration in the exploration–exploitation trade-off.

Two methodological weaknesses exist in the present study. First, SD conditions were not counterbalanced across task order. Second, baseline sleep patterns were not recorded or controlled for. With respect to the second point, the West Point cadets have a highly regimented lifestyle and likely had a very homogeneous sleep pattern. Additionally, their previous night's sleep was peer monitored and maintained to reflect eight hours of time in bed. Future studies will need to expand and replicate these findings while counterbalancing run order and controlling baseline sleep levels.

Recently, it has become clear that sleep deprivation does not have a straightforward effect on reasoning and decision-making. With this experiment, it appears the story is more complicated still. Subsequent investigation should work to further understand the time course, magnitude and specificity of sleep deprivation in the exploration–exploitation trade-off. Furthermore, the role of sleep consolidation should be examined in the task (Walker et al. 2003). The current study has demonstrated that decision making under sleep-deprivation is altered in a complex way that interacts with the explorative and exploitative properties of the task at hand.

Acknowledgements

The authors are pleased to acknowledge the support of the Army, grant #W911NF-07-2-0023, through The Center for Strategic and Innovative Technologies at The University of Texas at Austin for the research described here in this report.

References

- Belenky G, Penetar DM, Thorne D. 1994. Food components to enhance performance. Washington (DC): National Academy Press. The effects of sleep deprivation on performance during continuous combat operations. p. 127–135.
- Chee MW, Tan JC, Zheng H. 2008. Lapsing during sleep deprivation is associated with distributed changes in brain activation. *J Neurosci.* 28:5519–5528.
- Chuah YML, Venkatraman V, Dinges DF, Chee MWL. 2006. The neural basis of interindividual variability in inhibitory efficiency after sleep deprivation. *J Neurosci.* 26:7156–7162.
- Daw ND, O’Doherty JP, Dayan P, Seymour B, Dolan RJ. 2006. Cortical substrates for exploratory decisions in humans. *Nature.* 441:876–879.
- Drummond SPA, Gillin JC, Brown GG. 2001. Increased cerebral response during divided attention task following sleep deprivation. *J Sleep Res.* 10:85–92.
- Drummond SPA, Paulus MP, Tapert SF. 2006. Effects of two nights sleep deprivation and two nights recovery sleep on response inhibition. *J Sleep Res.* 15:261–265.
- Harrison Y, Horne JA. 1998. Sleep loss affects risk-taking. *J Sleep Res.* 7(Supplement 2):113.
- Harrison Y, Horne JA. 2000. The impact of sleep deprivation on decision making: A review. *J Exp Psychol Appl.* 6:236–249.
- Herscovitch J, Stuss D, Broughton R. 1980. Changes in cognitive processing following short-term cumulative partial sleep deprivation and recovery oversleeping. *J Clin Exp Neuropsychol.* 2:301–319.
- Ishii S, Yashida W, Yoshimoto J. 2002. Control of exploitation-exploration meta-parameter in reinforcement learning. *Neural Netw.* 15:665–687.
- Maddox WT, Glass BD, Wolosin SM, Savarie ZR, Bowen C, Matthews MD, Schnyer DM. 2009. The effects of sleep deprivation on information-integration categorization performance. *Sleep.* 32:1439–1448.
- McKenna BS, Dickinson DL, Orff HJ, Drummond SP. 2007. The effects of one night of sleep deprivation on known-risk and ambiguous-risk decisions. *J Sleep Res.* 16:245–252.
- Neri DF, Shappell SA, DeJohn CA. 1992. Simulated sustained flight operations and performance, part 1: effects of fatigue. *Mil Psychol.* 4:137–155.
- Ratcliff R, Van Dongen HPA. 2009. Sleep deprivation affects multiple distinct cognitive processes. *Psychon Bull Rev.* 16:742–751.
- Roberts ME, Goldstone RL. 2006. EPICURE: Spatial and knowledge limitations in group foraging. *Adapt. Behav.* 14:291.
- Strangman G, Thompson JH, Strauss MM, Marshburn TH, Sutton JP. 2005. Functional brain imaging of a complex navigation task following one night of total sleep deprivation: a preliminary study. *J Sleep Res.* 14:369–375.
- Tomasi D, Wang RL, Telang F. 2008. Impairment of attentional networks after one night of sleep deprivation. *Cereb Cortex.* 18:1.

- Venkatraman V, Chuah YML, Huettel SA, Chee MWL. 2007. Sleep deprivation elevates expectation of gains and attenuates response to losses following risky decisions. *Sleep*. 30:603–609.
- Walker MP, Brakefield T, Hobson AJ, Stickgold R. 2003. Dissociable stages of human memory consolidation and reconsolidation. *Nature*. 425:616–620.
- Williamson AM, Feyer A, Mattick RP, Friswell R, Finlay-Brown S. 2000. Developing measures of fatigue using an alcohol comparison to validate the effects of fatigue on performance. *Accid Anal Prev*. 33:313–326.
- Worthy DA, Maddox WT, Markman AB. 2007. Regulatory fit effects in a choice task. *Psychon. Bull Rev*. 14:1125–1132.