



Bayesian Versus Heuristic-Based Choice Under Sleep Restriction And Suboptimal Times Of Day

By: **David L. Dickinson** and Todd McElroy

Abstract

This paper examines the impact of a commonly experienced adverse cognitive state on decision making under uncertainty. Specifically, we administer an at-home sleep restriction protocol combined with random assignment to the time-of-day for decision making. Thus, we induce sleepiness in our subjects via sleep restriction as well as suboptimal time-of-day prior to administration of a Bayesian choice task. The specific task used discriminates between Bayesian choices that coincide with more simple reinforcement heuristic choices (in "Easy" trials) versus those that do not (in "Hard" trials), which is ideal given our underlying hypothesis that sleepy subjects are more likely to use simple heuristics. We first show that both circadian mismatch and sleep restriction significantly increase subjective sleepiness – this documents protocol validity. Our key behavioral results are that sleepy subjects are more likely to make a Bayesian inaccurate decision and more likely to make decisions consistent with a simple reinforcement heuristic, particularly in more cognitively difficult "Hard" trials. Secondary results show that stimulation of subject affect increased used of the simple decision heuristic but, when combined with sleep restriction, increased affect may increase task motivation and improve choice accuracy. These results offer new insights into the likely impact of sleepiness on decision making under uncertainty and highlight the potential negative impact on such cognitive states may have on accurate formation of probability assessments.

Dickinson, David L. and McElroy, Todd (2017). Bayesian Versus Heuristic-Based Choice Under Sleep Restriction and Suboptimal Times of Day. IZA Discussion Paper No. 10985. Publisher version of record available at: <https://ssrn.com/abstract=3037866>

DISCUSSION PAPER SERIES

IZA DP No. 10985

**Bayesian versus Heuristic-Based Choice
under Sleep Restriction and Suboptimal
Times of Day**

David L. Dickinson
Todd McElroy

SEPTEMBER 2017

DISCUSSION PAPER SERIES

IZA DP No. 10985

Bayesian versus Heuristic-Based Choice under Sleep Restriction and Suboptimal Times of Day

David L. Dickinson

Appalachian State University, IZA and ESI

Todd McElroy

Florida Gulf Coast University

SEPTEMBER 2017

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Bayesian versus Heuristic-Based Choice under Sleep Restriction and Suboptimal Times of Day

This paper examines the impact of a commonly experienced adverse cognitive state on decision making under uncertainty. Specifically, we administer an at-home sleep restriction protocol combined with random assignment to the time-of-day for decision making. Thus, we induce sleepiness in our subjects via sleep restriction as well as suboptimal time-of-day prior to administration of a Bayesian choice task. The specific task used discriminates between Bayesian choices that coincide with more simple reinforcement heuristic choices (in “Easy” trials) versus those that do not (in “Hard” trials), which is ideal given our underlying hypothesis that sleepy subjects are more likely to use simple heuristics. We first show that both circadian mismatch and sleep restriction significantly increase subjective sleepiness – this documents protocol validity. Our key behavioral results are that sleepy subjects are more likely to make a Bayesian inaccurate decision and more likely to make decisions consistent with a simple reinforcement heuristic, particularly in more cognitively difficult “Hard” trials. Secondary results show that stimulation of subject affect increased used of the simple decision heuristic but, when combined with sleep restriction, increased affect may increase task motivation and improve choice accuracy. These results offer new insights into the likely impact of sleepiness on decision making under uncertainty and highlight the potential negative impact on such cognitive states may have on accurate formation of probability assessments.

JEL Classification: C91, D81, D91

Keywords: sleep restriction, sleep deprivation, reinforcement heuristic, Bayesian choice, experiments

Corresponding author:

David Dickinson
Walker College of Business
3090 Peacock Hall
Boone, North Carolina 28608-2037
USA

Email: dickinsondl@appstate.edu

INTRODUCTION:

According to recent data, chronic sleep restriction (i.e., ≤ 6 hr/night) is habitual for approximately 30% of U.S. adults (Schoenborn and Adams, 2010), and shift work of some sort is performed annually by over 20 million U.S. wage and salary workers (McMenamin, 2007). Insufficient sleep is also considered a public health problem by the U.S. Centers for Disease Control and Prevention, and so an investigation of how commonly experienced levels of sleep restriction and adverse circadian timing impacts decisions under uncertainty is timely. Specifically, a main objective of ours is to examine the impact of sleepiness on Bayesian versus heuristic-based choices using a protocol that has high ecological validity and yet still constitutes a proper experiment (i.e., random assignments and sleep treatment manipulations).

Previous studies on sleep and Bayesian choice have used either observational (non-manipulated) sleep levels, or in-lab total sleep deprivation protocols. Our experimental protocol has high ecological validity in that we focus on levels of mild but chronic sleep restriction and suboptimal times-of-day that are highly relevant to real-world decision makers—the levels of sleep restriction and circadian suboptimal timing we study have direct relevance to tens of millions of adults (in the U.S. alone). Additionally, because our protocol allows individuals to undergo the experimental sleep manipulations in their natural environment with no real restrictions on compensatory behaviors, our primary data generation has high external validity.

Choice under uncertainty is fundamental to many environments and therefore motivates our use of a Bayesian choice paradigm. Given the novelty of our sleep protocol, we chose a vetted Bayesian task from the existing literature (Charness and Levin, 2005). The task discriminates between “Easy” trials, where Bayesian choices are empirically indistinguishable from simple reinforcement heuristic choices, and “Hard” trials, where the reinforcement heuristic produces Bayesian error. This task is useful for our purposes given our underlying hypothesis (articulated later) that sleep restricted or suboptimal time-of-day decisions are more likely to rely on less deliberative cognitive processes, such as the use of heuristics.

BACKGROUND:

Attitudes towards risk are an important component of decision making under risk as well as uncertainty. Regarding risky choice, the literature has documented tendencies to increase risk taking when sleep deprived (e.g., Killgore et al, 2012), although others have found results consistent with desensitization to risk following sleep deprivation in a more pure risk taking task (McKenna et al, 2007).¹ These previous studies typically involve extreme levels of laboratory sleep deprivation in a hyper-controlled environment that may not reflect a typical sleep

¹ We note, however, than many studies in the sleep literature administer tasks that are not well-designed to disentangle risky choice versus uncertainty (e.g., ambiguous gambles) choices, or choices over gains versus loss domains. For example, the Iowa Gambling task is well-known in sleep research as a way to evaluate risky choice—the Killgore et al, 2012, among many others, uses this task in sleep research. However, the task does not evaluate pure risk, but rather ambiguity, and it allows for gains or losses from a given decision, thus confounding choices over the gains and loss domains. The McKenna et al. (2007) study was designed to address these concerns.

deprivation experience. Such studies, while valuable and often the necessary first step in answering a basic research question, trade high levels of internal control for external validity. Alternatively, research using observational sleep levels benefits from the field-relevance of voluntary sleep choice, but such research does not offer the advantages of experimental manipulation or random assignment of sleep levels (see Dickinson et al, 2015, for a specific example of voluntary sleep choice on Bayesian decisions). Somewhat less common are studies looking at the impact of suboptimal circadian timing of a decision and its effects on risky choice. Castillo et al. (2017) is a recent study that found increased tendency to choose riskier asset bundles at suboptimal times of day that were similar to those we study in the present paper.

Because our focus is on decision-making under uncertainty, we highlight two particular studies most related to this paper. Dickinson and Drummond (2008) and Dickinson et al (2015) examine the impact of sleep restriction on outcomes in a Bayesian task environment. The task paradigm differs from the present paper, but consistent results are reported across the two studies. One study administered a laboratory imposed total sleep deprivation protocol (22-24 cumulative wake hours: Dickinson and Drummond, 2008) while the other studied objective measures of voluntary sleep choice and focused on those with ≤ 6 hours of sleep per night over the one week prior to decision making (Dickinson et al, 2015). Both papers found that restricted sleep led one to place relatively less decision weight on new evidence relative to existing information in the Bayesian task used. These studies are informative for the present paper because they highlight the fact that sleepy subjects appear less likely to engage high-level cognitive processes necessary to incorporate multiple sources of information into a decision.² Horne (1993) notes a general connection between sleep deprivation and decreased prefrontal activation. Although the recent literature is a bit more mixed, results general show harm to decision quality even if this is associated with increases in prefrontal activation of certain areas.³ Our hypotheses are therefore derived from the existing literature that finds results consistent with reduced use of beneficial deliberation when sleep restricted or at a lower point of circadian alertness.

The present study administers a unique at-home sleep manipulation protocol with random assignment of time-of-day for decisions to evaluate decisions in the Bayesian switching task of Charness and Levin (2005). The specifics of the task are discussed in the next section, but the important feature is the distinction between “Easy” versus “Hard” trials noted in the Introduction above. In Easy trials, a simple reinforcement heuristic can be employed to make decisions that are Bayesian-accurate. Because of this, a decision maker that uses the simple decision heuristic may paradoxically make more accurate decision assessments if deliberative thinking poses any risk of overthinking the decision.⁴ Hard trials separate what a heuristic decision-maker would choose relative to a Bayesian decision-maker, thus allowing one to separate and identify

² Poudel et al (2017) document the neural correlates of well-rested decision makers using the specific Bayesian task of these prior studies.

³ See Dickinson and McElroy (in press) for a concise summary of the recent literature on prefrontal activation in high-level decision making following sleep loss.

⁴ For example, Achtziger and Alós-Ferrer (2014) develop a 2-stage model of decision making that posits increased likelihood of error is one chooses a deliberative cognitive system when a more automatic would be appropriate for the choice.

individuals based on likely cognitive approach to the decision. Because of the 2-stage nature of each trial in the Bayesian task, a final consideration worth noting is the reward anticipation experienced by a subject who is shown a black ball in Stage 1—a black ball is the “payoff ball” for the Stage 2 decision that follows. The positive emotion or reward anticipation noted by Charness and Levin (2005) is relevant because sleep deprivation has been shown to enhance reward anticipation (see Venkatraman, 2007, 2009, and 2011). These studies offer an example of how neural activation changes following sleep deprivation are not beneficial to decision quality. In this case, the reason is the introduction of an optimism bias where the individual has heightened anticipation of the monetary reward. In the context of the present study, this implies sleep restricted or circadian mismatched subjects may experience an even greater increase in positive affect than usual after a Stage 1 black ball draw.

EXPERIMENTAL DESIGN:

Here we present the essential features of the sleep manipulations, and direct the reader elsewhere for a detailed examination that includes analysis of sample selection, attrition, and validity of the at-home sleep protocol (Dickinson et al., 2017). Specifically, our unique mixed design combines experimentally manipulated sleep levels with random assignment to a more or less preferred time-of-day for the decision sessions. The at-home sleep protocol is valid in that we have documented significantly increased sleepiness when subjects are either sleep restricted or at a sub-optimal time of day compared to when well-rested or at a more preferred time-of-day (Dickinson et al., 2017)

A short online survey was administered every semester for several years, and within that survey we included a validated instrument for measuring morningness/eveningness preferences in order to identify validated morning-types and evening-types within a regularly maintained database. Then, we randomly assigned morning-types and evening-types, *ex ante*, to a Morning (7:30am-9:00am) or Evening (10:00pm-11:30pm) experiment group prior to sending an email invitation for the main study. We refer to subjects assigned a more preferred time of day as “circadian matched” and those at less preferred times of day as “circadian mismatched.” This circadian match/mismatch dimension of the sleep protocol effectively generates random assignment to a more or less optimal time-of-day for decision making.

The recruitment email for the main study invited morning-type and evening-type subjects to participate in a 3-week study requiring adherence to a prescribed sleep schedule and 3 in-lab sessions. Each subject was asked to spend 1 week sleeping 5-6 hr/night (sleep-restricted: SR) and 1 week sleeping 8-9 hr/night (well-rested: WR). These treatment weeks constitute weeks 1 and 3 of the protocol, with the treatment order counterbalanced across groups. For all subjects, week 2 was an *ad lib* sleep week to washout the effect of week 1 prior to commencing the week 3 schedule. Compliance was assessed using validated wrist-worn actigraphy devices common to clinical and sleep research. As such, nightly sleep levels were measured passively but objectively using validated instrumentation. The sleep restriction dimension of the design is therefore within-subjects, whereas the circadian mismatch is a between-subjects manipulation.

Importantly, the Bayesian task was administered within both decision sessions that occurred at the end of weeks 1 and 3. Thus, each subject made choices in the Bayesian task under both SR and WR conditions (at either a constant optimal or suboptimal time of day).

Subjects were compensated a fixed \$80 for compliance with the sleep protocol details, wearing the actigraphy device, and providing completed sleep diaries at the end of the 3 weeks. This fixed payment occurred several days after completion of the study to give researchers time to download sleep data and verify good-faith effort to comply with the sleep prescription.⁵ Compensation for outcomes in the Bayesian decision task was separate from this fixed compensation, and it occurred in cash at the end of a decision task session.

The Bayesian Decision Environment

We administered a modified version of the Bayesian choice task in Charness and Levin (2005). Figure 1 described the task, in which a trial or round of the task involves two stages. For each trial of the task, subjects see the 2x2 stimulus showing varied numbers of black and white balls in each cell of a matrix. Subject only make a decision in Stage 2, and the subject is paid \$10 if a black ball is drawn in stage 2 of a randomly selected trial from this task (\$0 otherwise) and therefore has a payoff interest in making a decision to maximize the probability of a black ball being drawn in Stage 2.⁶ In Stage 1, Nature selects the TOP or BOTTOM row, which will be the row used for both stages of that trial, with 50% chance for each (described to subjects as flip of a fair coin). Subjects are not, however, made aware of the outcome of Nature's coin flip. Rather, in Stage 1 of the trial subjects are told that one ball will be drawn with replacement from either the LEFT or the RIGHT column of the task matrix. Subjects observe the outcome of that draw (e.g., Stage 1 column is RIGHT and the resultant draw shown to the subject was a white ball). Thus, Stage 1 provides information regarding the row Nature had drawn for that trial. More specifically, if subjects are told the Stage 1 column is RIGHT, and the ball drawn is black, then this eliminates all uncertainty—the subject should then know that the row is UP for this trial. In Stage 2 of the trial, subjects are then asked to choose the column from which a ball will be drawn with the payoff implications noted above.

The example above described an Easy trial, because a ball drawn in stage 1 from RIGHT tells the subject the row with certainty. If a black ball is drawn in Stage 1, then a subject maximizes her expected payoff by selecting RIGHT for the Stage 2 draw. If, however, a white ball is drawn in Stage 1 from RIGHT, then the subject maximizes her expected payoff by choosing to have the Stage 2 ball drawn from LEFT. Importantly, in these Easy trials (where Stage 1 draw is from

⁵ It is worth noting that we used a fairly loose standard of “compliance” for the choice to release full payment of the fixed compensation (\$80) for study participation. That standard was different than what we considered “compliant” data for purposes of the analysis of the decision data discussed in our Results section. For data analysis, we considered a subject “compliant” if there was an objectively measured difference in nightly sleep between the WR and SR weeks of at least 60 minutes per night for that subject (see Dickinson et al, 2017, for additional discussion of compliance standard choice).

⁶The reader may note that this is a slight modification of the Charness and Levin (2005) environment. Our modifications focus attention on the Stage 2 choice for our purposes, and we modify the number of black and white balls in each cell of other task matrix to make the difference in expected payoffs between an accurate versus inaccurate (Bayesian) choice a bit larger.

LEFT) the payoff maximizing or Bayesian optimal choice, perfectly coincides with the choice a subject would make if following a simple “win-stay, lose-switch” reinforcement heuristic. Using this heuristic, if a subject sees a black ball drawn in Stage 1 then, because a black ball is what would give the subject a positive payoff, the subject will “stay” and choose the same column for Stage 2 as was used for Stage 1. If, however, a white ball is drawn from LEFT in Stage 1, then the heuristic would lead a subject to switch and choose RIGHT for the Stage 2 draw. As Charness and Levin (2005) note, one cannot empirically distinguish a Bayesian subject from a subject who uses the simpler reinforcement heuristic in Easy trials.

Hard trials, on the other hand, are those where the Stage 1 draw is from the LEFT column. Here, a Bayesian subject would maximize her expected payoff by switching to RIGHT for Stage 2 if the computer draws a black ball from LEFT in Stage 1 (i.e., the black ball reveals it is more likely one is in the UP row for that trial). The reinforcement-heuristic, however, would cause one to “stay” and have the Stage 2 draw from the LEFT column even though this does not maximize expected payoff. Likewise, a white ball from LEFT in Stage 1 would cause a reinforcement-heuristic subject to switch to RIGHT for Stage 2, even though the Bayesian choice is to stay with the LEFT column for Stage 2. In total, subjects are administered a set of 40 trials: 20 Easy and 20 Hard trials randomly ordered across the 40 trials.

Hypotheses

As noted, the Hard trials separate a Bayesian from a reinforcement-heuristic decision making. Thus, if sleepiness promotes less deliberative and more automatic or heuristic-based decisions, we predict reduced accuracy in Hard trials. In a strict sense, we predict no difference in accuracy on Easy trials, since Stage 1 eliminates uncertainty. However, if one assumes it is possible to over-deliberate, then we may hypothesize increased Bayesian accuracy on Easy trials when sleep restricted or circadian mismatched—here, the increased use of a simple reinforcement heuristic actually works in favor of accuracy for a sleepy subjects since reinforcement choices will be Bayesian, and a subjects attempting to deliberate when sleepy may be more prone to decision error. The implicit assumption that over-deliberation is possible, which may be questionable to some, is somewhat consistent with Achtziger and Alós-Ferrer (2014), who assume that the decision to use a deliberative cognitive process increases the potential for error if an automatic process is more appropriate. Our first hypotheses are therefore with respect to decision accuracy.

Hypothesis 1: Bayesian accuracy will be lower on Hard trials (i.e., when reinforcement choices diverge from Bayesian choices) when sleep restricted or circadian mismatched compared to when well-rested or at a preferred time-of-day.

Hypothesis 2: Bayesian accuracy will be higher on Easy trials (i.e., when reinforcement choices coincide with Bayesian choices) when sleep restricted or circadian mismatched compared to when well-rested or at a preferred time-of-day.

The task allows for evaluation of emotion or “affect” on decisions, which we seek to exploit in our data analysis. Charness and Levin (2005) identify that the Stage 1 draw likely generates positive affect if a black ball is drawn. Their assumption is that this positive affect then leads to more emotion-based choices in Stage 2, which favors the reinforcement heuristic. In our simplified version of this Bayesian task we do not compensate subjects if a black ball is drawn in Stage 1 by the computer. Nevertheless, we hypothesized that a black ball in Stage 1 will generate a degree of positive affect that favors the reinforcement heuristic. Alternatively, we noted above that the positive affect from a black ball drawn in Stage 1 may enhance reward anticipation to a greater degree when sleepy. This would promote an increased propensity to use the reinforcement heuristic following a Stage 1 black ball draw for sleep restricted or circadian mismatched subjects. In other words, when a black ball is drawn in Stage 1 optimism may be primed more sleepy subjects. This mechanism leads to a hypothesis regarding how our sleep manipulations will change behavior differentially following a Stage 1 black ball draw.

Hypothesis 3: The reinforcement heuristic will be used more when a black ball is drawn in Stage 1 if sleep restricted or circadian mismatched compared to when well-rested or at a preferred time-of-day.

RESULTS:

Pooled analysis (collapsing across trials)

We report data on 119 subjects deemed compliant with the 3-week prescription (n=62 circadian matched; n=57 circadian mismatched). Analysis of data pooled across the set of 20 Easy and 20 Hard trials was first conducted on Bayesian choice by creating a Bayes Score $\in [0,20]$ for each trial type. Bayes Scores were significantly higher on Easy trials (16.95 ± 4.55) than on Hard trials (13.03 ± 4.67), and the difference is significant during both administrations of the task (Signed rank test, $p < .01$ in each instance). Random effects regressions (2 observations per subject) on Bayes Score are estimated with sleep restriction (SR) and circadian mismatch (MM) as indicator variables. An additional indicator variable for the second decision session is also included to account for learning or repeat administration effects for this task. The results are shown in Table 1.

The estimations indicate that subjects are Bayesian accurate on significantly more trials upon the second administration of the task—approximately one to two more trials correct out of the 20 total trials of Hard and Easy tasks, respectively. On Easy trials, MM condition leads to increased BayesScores, while on Hard trials the SR condition leads to significantly lower Bayes Scores. Robustness estimations are included in the Appendix and include alternative measures of scoring SR as a continuous variable or using self-reported sleepiness as replacement for the MM and SR indicator variables. The results are robust across specifications, with the exception of the estimations using sleepiness as a regressor. Here, whether using sleepiness ratings itself, or whether instrumenting sleepiness from sleep specific treatment and demographic variables, the additional estimations support Table 1 results only in the Hard trials. Thus, we consider it more

robust that subjects in the SR condition are significantly less Bayesian accurate on Hard trials, which support Hypothesis 1.⁷

Support for Hypothesis 1 but not Hypothesis 2 is perhaps not surprising given the statistically more powerful within-subjects SR condition compared to the between-subjects MM. It is also the case that the magnitude of the impact of SR on self-reported sleepiness is significantly higher than the magnitude of the MM impact (see Appendix Table A4 first-stage regression results). This likely results from the fact that the experimental manipulation of SR involve a full week of accumulated sleep restriction and thus is an overall more effective treatment manipulation.

Pooled trial analysis was also conducted on response times, though in a more exploratory fashion. Some have used response times as a valid indicator of cognitive process (see Kahneman, 2011), but we reserve judgment on the predicted effects of an adverse sleep state response times. For example, an implication of the response time theory in Achziger and Alós-Ferrer (2014) is that a sleepy subject may nevertheless produce systematically longer response times. This would be the case if sleepy subjects takes more time selecting whether to use an automatic or deliberative decision process (as distinguished from the typically longer response times that should be observed with deliberative compared to automatic decision processes once the process is selected). In other words, though sleepy subjects may be more likely to use a faster non-deliberative decision process, the resultant prediction on observable response times is unclear because a delay may exist in selecting/engaging the fast decision process when sleepy.

Response times (in milliseconds) were averaged across the set of 20 Hard and Easy trials for each subject and the results are also shown in Table 1. Here, average response times are longer on Hard trials ($790.04 \pm 306.97\text{ms}$) compared to Easy trials ($733.64 \pm 265.25\text{ms}$). However, the difference is only significant during the first administration of the task (Signed Rank test: $p < .01$) and average response times are statistically no different during the second administration of the task at the end of the 3-week experiment protocol. Table 1 shows that response times on both Easy and Hard trials are significantly quicker during the second task administration, and adverse sleep states of SR and MM both increase average response time. Given the known impact of sleep restriction on impaired executive function (i.e., generally impaired deliberative thought processes), the fact that sleepier subjects display long response times suggest caution in any interpretation of longer response times as simple indicators of increased deliberation. It seems clear that the interpretation of response time data is a more complex endeavor that would likely benefit greatly from complementary data to help accurately identify the decision process being used during decision making.

Reinforcement Heuristic estimations (trial-level analysis)

⁷As noted in the first-stage regression of the instrumental variables estimation approach in Appendix Table A4, the MM condition increases self-reported sleepiness significantly, but significantly less so than the SR condition. Additionally, we find in our sample that a subject in the compounded SR*MM condition actually report lower sleepiness levels during the decision session than a subject in either SR or MM by itself. This may result from subject engaging in countermeasures when both SR and MM. Because of this possibility, and the fact that self-reported sleepiness may not reflect objective sleep need (see Van Dongen et al, 2003), the principal trial-level analysis will use objective sleep condition as a regressor rather than self-reported sleepiness.

For the trial-level analysis, we focus on the key outcome measure at issue for the test of our hypotheses: consistency of choice with the reinforcement heuristic. We therefore score a dichotomous variable equal to one for each trial where a subject's choice is consistent with the reinforcement heuristic (and zero otherwise).⁸ Independent variables include the SR and MM treatments, the 2nd Decision Session indicator, an variable that indicates a Stage 1 black ball draw, a trial variable to capture within-session learning, and a gender indicator. We use random effects models to account for multiples observations per subject, and estimations include both a base model of reinforcement choice determinants (models (1) and (3)) that estimates only learning and main treatment effects, and models with additional covariates (models (2) and (4)). The expanded models include a term for gender, interaction terms between the sleep treatments (SR and MM) and the Stage 1 Black ball draw, and a Trial count variable. In Table 2 we report the marginal effects of each regressor on the probability of making a reinforcement choice.

Focusing first on Easy trials, we see in Table 2 that subject are more likely to make a reinforcement choice with each passing trial, as well as in the second decision session. Both of these are beneficial effects in the Easy trials where reinforcement choices are also Bayesian. We also find in model (1) that a black ball drawn in Stage 1 increases the probability that one makes a reinforcement heuristic choice, which is consistent with the hypothesis a Stage 1 black ball increases positive affect and promotes more automatic response. The estimated impacts of the sleep treatment variables, SR and MM, are statistically insignificant on Easy trials. In the expanded Easy trials model (2), we also estimate that females are 5% less likely to make a (beneficial) reinforcement choice. Regarding the impact of a Stage 1 black ball draw, we find that the increase in reinforcement-consistent choices occur only for subjects either in the SR or MM conditions. This offers support for Hypothesis 3 for Easy trials. We also find that MM subjects are more likely to use the reinforcement heuristic with each passing trial. In short, to the extent that SR or MM impact reinforcement choices, they increase the use of reinforcement, which is beneficial on Easy trials. This result helps explain the mechanism behind the Hypothesis 2 support in Table 1 regarding improved Easy trial accuracy for MM subjects.

The results for Hard trials, on the other hand, show notable differences. Repeat task or “learning” effects (*Trial* and *2nd Decision Session*) are in the direction of significant reductions in reinforcement heuristic choices on Hard trials. Thus, the consistency across across both Easy and Hard trial types is that Bayesian accuracy improves with repetition of the task. Female subjects are significantly more likely to make reinforcement choices in Hard trials, which is harmful for decision accuracy. The gender effect we report is therefore consistent across trial types and explains higher Bayesian error rates in this task among females.⁹ Also, the reinforcement priming of drawing a black ball in Stage 1 is significant but much larger in magnitude in Hard compared to Easy trials. Specifically, a black ball in Stage 1 increase the probability of making a (erroneous) reinforcement heuristic choice by over 20% in Hard trials (versus a small magnitude effect that was not robust across models (1) and (2)). Somewhat

⁸ Note that such a dichotomous variable is equivalent to an indicator of Bayesian choices in Easy trials but not-Bayesian choices in Hard trials.

⁹ Mann-Whitney tests of Bayes scores show females Bayes scores are significantly lower than male Bayes scores in both Easy trials ($p < .01$) and Hard trials ($p = .03$).

surprisingly, this effect of the Stage 1 black ball is *diminished* among subjects in the SR condition. So, our results do not support the general Hypothesis 3 claim that a black ball in Stage 1 will promote increase use of the reinforcement heuristic for SR or MM subjects. At least for SR subjects, the significant Stage 1 black ball effect promotes increased beneficial decision effects relative to being well-rested, perhaps due to increased motivation and focus.

A final analysis is conducted to explore differences in the types of error subjects make. More specifically, we distinguish between a “Stay Error” and a “Switch Error”. The variable Stay Error is coded as equal to one if a subject chose the same column as the Stage 1 draw, even though the Bayesian accurate choice would have to been to switch columns. Switch Error is equal to one if a subject switched columns though it would have been Bayesian to stay. Table 3 reports random effects probit regressions (marginal effects reported) separately for each error type and trial type. Results in Table 3 show that the improvements in accuracy due to learning involve reductions in both error types in Easy trials, although the magnitude of the learning effect is larger for reductions in Stay Errors. In Hard trials, however, learning across trials or sessions significantly reduces only Switch Errors. Again we find the significant impact regarding sleep manipulations only in the SR condition, where SR is estimated to increase Stay Errors on Easy trials but increase Switching Errors on Hard trials. The impact of SR on type of errors committed for a given trial type is in the same direction of the estimated impact of being female on error type committed relative to male subjects.

CONCLUSIONS:

Our results indicate that commonly experienced adverse sleep states have significant effects on decisions under uncertainty. We formulated three hypotheses regarding the impact of sleep restriction and circadian mismatch on outcomes in the Bayesian choice task administered. Our results support Hypothesis 1—sleep restriction significantly reduced Bayesian accuracy on Hard trials. There was weak support for Hypothesis 2, where we estimated a weak but positive effect of circadian mismatch on Bayesian accuracy in Easy trials. The robustness analysis identifies the SR effect as stable across numerous specifications, while support for Hypothesis 2 was not robust. This may be due to the statistically and practically more powerful impact of a within-subjects SR manipulation whereby sleep debt is accumulated across an entire week (as opposed to a single night circadian mismatch for the decision session).

Tests of Hypotheses 1 and 2 are also indirectly embedded in our analysis of the probability of making a choice consistent with the reinforcement heuristic. Table 2 results showed that when reward anticipation is primed through the draw of a black (payoff) ball in Stage 1 of the Easy trials, then both SR and MM subjects increased their use of the reinforcement heuristic. This supports Hypothesis 3 for Easy trials and also offers qualified support for Hypothesis 2 (i.e., the hypothesis is supported when affect is primed in Stage 1 because increased reinforcement heuristic use improves accuracy on Easy trials). Table 2 analysis also supports Hypothesis 1 for Hard trials (SR is shown to increase reinforcement choice in Hard trials where it harms accuracy), but does not support Hypothesis 3 regarding Hard trials. In sum, we report a robust

result of SR harming Bayesian accuracy on Hard trials due to increased use of the reinforcement heuristic. The mixed results regarding Hypothesis 3 are not supportive of our initial conjecture regarding the impact of a Stage 1 Black Ball and sleepiness, but the results may be consistent with a different hypothesis not initially considered. Namely, the Stage 1 black ball draw may heighten task motivation by drawing attention to the monetary incentives at stake in the decision task. While our ex ante view was that this would engage detrimental overconfidence or optimism regarding the Stage 2 outcomes, the specifics of our task are not particularly conducive towards an optimism bias on Stage 2 choice. Rather, this result may highlight how at least mild adverse sleep states may focus one's effort. Of course, this is post hoc reasoning, but may offer insights to consider. It has been shown in the sleep literature that incentivized tasks may help combat some of the detrimental task performance effects of sleep deprivation (see Alhola and Polo-Kantola, 2007, and sources within), and so this result may be worth exploring more systematically in future research.

While we did not have a formal hypothesis for the types of errors subjects would commit under adverse sleep conditions, these results may also be of interest for further investigation. Specifically, it may be the case that choosing from the same column in Stage 2 as was used for Stage 1 is a default or status quo option. This may then imply an interesting interpretation of the sleep restriction effects we reported. Namely, it may be that SR causes more errors of omission (Stay Errors) on Easy trials, but more errors of commission (Switch Errors) on Hard trials.

Our main finding is an important one and adds to existing results in the literature on sleep and Bayesian choice. When the task environment is challenging and clearly separates Bayesian from reinforcement decision makers, subjects under conditions of mild but commonly experienced sleep restriction are more likely than rested subjects to make inaccurate assessments. Previous findings reported a decreased decision weight placed on important new information when sleep deprived (Dickinson and Drummond, 2008) or voluntary low sleep levels (Dickinson et al, 2015), but these earlier studies used a distinct protocol and found no evidence that these significant decision model effects impacted Bayesian accuracy. In the present study, we find a significant impact on accuracy when sleep restricted. This finding has important implications in any environment where individuals must incorporate new evidence into existing information to formulate beliefs. For example, health decisions represent one such environment where existing information (base rate risks) may not coincide with new evidence (e.g., a positive test result) and may therefore present a difficult Bayesian assessment environment. While sleepy decision makers may yet make accurate assessments if the decision environment happens to be "easy", the accurate incorporation of new information is the preferred approach to make accurate assessments and will be necessary for accurate choice in more challenging decision environments.

REFERENCES

Achtziger, A., & Alós-Ferrer, C. (2014). "Fast or rational? A response-times study of Bayesian updating." *Management Science*, 60: 923–993.

- Alhola P, Polo-Kantola P (2007) "Sleep deprivation: impact on cognitive performance." *Neuropsychiatric Disease and Treatment*. 3(5): 553-567.
- Castillo, Marco, David L. Dickinson, and Ragan Petrie (2017) "Sleepiness, choice consistency, and risk preferences." *Theory and Decision*, 82(1): 41-73.
- Charness, G., and D. Levin. 2005. "When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect." *American Economic Review*, 95(4): 1300-1309.
- Dickinson, David L., and Sean P.A. Drummond (2008) "The Effects of Total Sleep Deprivation on Bayesian Updating." *Judgment and Decision Making*, 3(2): 181-90.
- Dickinson, David L., Sean P. A. Drummond, and Jeffrey Dyche (2016) "Voluntary sleep choice and its effects on Bayesian decisions." *Behavioral Sleep Medicine*, 14(5): 501-513
- Dickinson, David L, Sean P. A. Drummond, Todd McElroy (2017) "The viability of an ecologically valid chronic sleep restriction and circadian timing protocol: An examination of sample attrition, compliance, and effectiveness at impacting sleepiness and mood." *PLoS ONE*, 12(3): e0174367.
- Dickinson, David L., and Todd McElroy (2017) "Sleep restriction and circadian effects on social decisions." *European Economic Review*, 97(August): 57-71.
- Horne, J.A. (1993) "Human sleep, sleep loss and behavior. Implications for the prefrontal cortex and psychiatric behavior." *Brit. J. Psychiat.* 162: 413-419.
- Kahneman, D. *Thinking, fast and slow*. Macmillan, 2011.
- Killgore, W.D., N.L. Grugle, and T.J. Balkin. (2012) "Gambling when sleep deprived: Don't bet on stimulants." *Chronobiology International*, 29: 43-54.
- McKenna, Benjamin., David L. Dickinson, Henry Orff, and Sean P.A. Drummond (2007) "The Effects of One-Night of Sleep Deprivation on Known-Risk and Ambiguous-Risk Decisions." *Journal of Sleep Research*, 16(3): 245-52.
- McMenamin, T.M. 2007. "A time to work: Recent trends in shift work and flexible schedules." *Monthly Labor Review*, Dec.: 3-15.
- Poudel, Govinda, Anjan Bhattarai, David L. Dickinson, Sean P.A. Drummond (2017) "Neural correlates of decision making during a Bayesian choice task." *NeuroReport*, 28(4): 193-199.
- Schoenborn, C.A., and P.F. Adams. 2010. "Health behaviors of adults: United States, 2005-2007. National Center for Health Statistics. *Vital Health Stat, series 10, Data from the National Health Survey* 245: 1-132.

- Van Dongen, H.P.A., G. Maislin, J.M. Mullington, and D.F. Dinges (2003). “The cumulative cost of additional wakefulness: dose-response effects on neurobehavioral functions and sleep physiology from chronic sleep restriction and total sleep deprivation.” *Sleep*, 26(2): 117-1269.
- Venkatraman V, Chuah L, Huettel SA, Chee MW (2007). “Sleep deprivation elevates expectation of gains and attenuates response to losses following risky decisions.” *Sleep*. 30: 603-609.
- Venkatraman V, Huettel SA, Chuah LY, Payne JW, Chee MW (2011) “Sleep deprivation biases the neural mechanisms underlying economic preferences.” *Journal of Neuroscience*. 31(10): 3712-3718.
- Venkatraman V, Payne JW, Bettman JR, Luce MF, Huettel SA (2009) “Separate neural mechanisms underlie choices and strategic preferences in risky decision making.” *Neuron*. 62(4): 593-602.

FIGURE 1: The Bayes v. Reinforcement task

A coin flip picks UP or DOWN at the beginning of Stage 1 of each choice round (50% chance each for UP or DOWN). After the unknown row choice, LEFT or RIGHT is chosen by computer to complete Stage 1, with a payoff incentive to subject if a black ball is drawn. For Stage 2, the unknown row choice remains fixed, subjects know this, but subjects now choose whether or not to pick the same column as was picked for them in Stage 1.




Table 1: Treatment impacts on Response Times and Bayes Scores

(data pooled across the 20 trials of Easy or Hard choices for a given subject for each Session)
Coefficient estimates shown (standard errors in parenthesis)

N=238 (119 subjects)	Response Time estimations		Bayes Scores Estimations	
	Dep Var Average RT (ms) Easy Trials	Dep Var Average RT (ms) Hard Trials	Dep Var Bayes Score Easy Trials	Dep Var Bayes Score Hard Trials
Constant	731.55 (33.04)***	809.13 (37.52)***	15.46 (.57)***	12.66 (.60)***
2 nd Decision Session	-117.65 (.21.87)***	-177.55 (24.99)***	1.89 (.39)***	1.02 (.44)**
SD	44.14 (21.87)**	64.21 (24.99)***	-.15 (.39)	-.96 (.44)**
MM	81.10 (41.53)*	78.47 (47.08)*	1.28 (.72)*	.71 (.73)
Chi-squared	39.63***	64.59***	27.72***	12.58***

Note: Response times (in milliseconds) are averaged over the set of trials (Easy vs Hard trials separated). Bayes scores are the total number of Bayesian correct responses out of 20 for the set of trials (Easy vs Hard trials separated). Two observations on each of the 119 subjects exist (1st and 2nd experiment session) and estimation include random effects error structure at the subject level.

*, **, *** indicate significance at the .10, .05, and .01 level, respectively, for the 2-tailed test.

Table 2: Use of the Reinforcement Heuristic (at individual trial level)

Marginal effects shown (standard errors in parenthesis)

	Dependent Variable: Use of Reinforcement Heuristic (=1)			
	Easy Trials (n=4758)		Hard Trials (n=4758)	
	(1)	(2)	(3)	(4)
<i>Trial</i>	.0012 (.0003)***	.0007 (.0004)	-.0020(.0006)***	-.0005 (.0010)
<i>2nd Decision Session</i>	.0680 (.0141)***	.0670 (.0140)***	-.0570 (.0139)***	-.0539 (.0138)***
<i>Stage 1 Black Ball</i>	.0205 (.0071)***	-.0006 (.0103)	.2124 (.0140)***	.2541 (.0240)***
<i>SD</i>	-.0018 (.0062)	-.0116 (.0109)	.0533 (.0139)***	.1025 (.0250)***
<i>MM</i>	.0354 (.0272)	-.0055 (.0286)	-.0395 (.0414)	.0208 (.0497)
<i>female</i>	---	-.0512 (.0301)*	---	.0825 (.0420)**
<i>S1BB*SD</i>	---	.0237 (.0130)*	---	-.0620 (.0278)**
<i>S1BB*MM</i>	---	.0217 (.0130)*	---	-.0226 (.0280)
<i>SD*Trial</i>	---	.0000 (.0000)	---	-.0000 (.0000)
<i>MM*Trial</i>	---	.0013 (.0006)**	---	-.0019 (.0012)
Chi-squared	155.29***	165.82***	266.59***	277.22***

Note: Results from random effect probit regressions (random effects at subject level). Standard errors are calculated using the Delta-method.

*, **, *** indicate significance at the .10, .05, and .01 level, respectively, for the 2-tailed test.

Table 3: Use of the Reinforcement Heuristic (at individual trial level)

Marginal effects shown (standard errors in parenthesis)

	Dependent Variables: Stay Error=1 implies if subject stayed but should have switched to other column in Stage 2. Switch Error=1 if subject switched but should have chosen the same column in Stage 2.			
	Easy Trials (n=4760)		Hard Trials (n=4760)	
	Stay Error=1	Switch Error=1	Stay Error=1	Switch Error=1
<i>Trial</i>	-.0006 (.0002)***	-.0002 (.0001)**	-.0003 (.0005)	-.0013 (.0004)***
<i>2nd Decision Session</i>	-.0285 (.0069)***	-.0150 (.0058)***	-.0103 (.0118)	-.0418 (.0089)***
<i>female</i>	.0299 (.0163)*	.0064 (.0080)	.0262 (.0309)	.0487 (.0206)**
<i>SD</i>	.0080 (.0048)*	-.0014 (.0020)	.0180 (.0118)	.0248 (.0086)***
<i>MM</i>	-.0089 (.0141)	-.0102 (.0087)	-.0259 (.0299)	.0020 (.0193)
Chi-squared	57.87***	70.72***	5.35	55.35***

Note: Results from random effect probit regressions (random effects at subject level). Standard errors are calculated using the Delta-method.

*, **, *** indicate significance at the .10, .05, and .01 level, respectively, for the 2-tailed test.

APPENDIX: Robustness analysis of pooled (by trial) outcomes

Tables A1 and A2 shows results from estimations analogous to Table 1 in the main manuscript, but with alternative scoring of sleep condition. Specifically, continuous variables are introduced that measures one’s objectively measured nightly sleep level the week prior to decisions, *Sleep Quantity*, or that same objective measure subtracted from one’s perceived nightly sleep need, *Personal SD*. Either measure is intended to capture degree of sleep restriction intensity. However, it should be noted that factors not captured in our data set (other than the sleep treatment manipulation) itself may be responsible for the differences in those with difference levels of nightly sleep or sleep restriction in either week. Thus, the dichotomous scoring of SD in the main text is the most clean measure of the treatment manipulation for econometric estimation.

Table A1: Treatment impacts on Response Times and Bayes Scores

(data pooled across the 20 trials of Easy or Hard choices for a given subject for each Session)

Coefficient estimates shown (standard errors in parenthesis)

N=238 (119 subjects)	Response Time estimations		Bayes Scores Estimations	
	Dep Var Average RT (ms) Easy Trials	Dep Var Average RT (ms) Hard Trials	Dep Var Bayes Score Easy Trials	Dep Var Bayes Score Hard Trials
Constant	737.56 (37.11)***	799.91 (42.23)***	15.27 (.64)***	12.86 (.68)***
2 nd Decision Session	-120.42 (22.11)***	-179.59 (25.12)***	1.92 (.39)***	1.04 (.44)**
Personal SD	.16 (.18)	.38 (.20)*	.001 (.003)	-.006 (.003)*
MM	84.25 (41.77)**	86.11 (47.57)*	1.30 (.72)*	.58 (.73)
Chi-squared	35.62***	60.98***	27.60***	11.05**

Note: *Personal SD* is the difference between the objectively measured average nightly sleep level (in min/night) for the week preceding the decision and the perceived sleep need (in min/night) of that subject (elicited during the online sleep survey at a separate previous point in time)—that is, a personalized sleep deprivation measure. Only compliant subjects are included, but variation in *Personal SD* reflects different degrees of sleep restriction or rest in subjects that are handicapped for the subjects subjective sleep need. Response times (in milliseconds) are averaged over the set of trials (Easy vs Hard trials separated). Bayes scores are the total number of Bayesian correct responses out of 20 for the set of trials (Easy vs Hard trials separated). Two observations on each of the 119 subjects exist (1st and 2nd experiment session) and estimation include random effects error structure at the subject level.

*, **, *** indicate significance at the .10, .05, and .01 level, respectively, for the 2-tailed test.

Table A2: Treatment impacts on Response Times and Bayes Scores

(data pooled across the 20 trials of Easy or Hard choices for a given subject for each Session)

Coefficient estimates shown (standard errors in parenthesis)

N=238 (119 subjects)	Response Time estimations		Bayes Scores Estimations	
	Dep Var Average RT (ms) Easy Trials	Dep Var Average RT (ms) Hard Trials	Dep Var Bayes Score Easy Trials	Dep Var Bayes Score Hard Trials
Constant	898.86 (78.37)***	1040.59 (89.63)***	14.70 (1.39)***	9.60 (1.54)***
<i>2nd Decision Session</i>	-117.50 (21.85)***	-177.74 (.25.03)***	1.88 (.39)***	1.04 (.44)**
<i>Sleep Quantity</i>	-.38 (.19)**	-.52 (.22)**	.002 (.003)	.007 (.004)*
<i>MM</i>	82.18 (41.64)**	79.96 (47.23)*	1.28 (.72)*	.69 (.73)
Chi-squared	39.57***	63.50***	27.88***	10.73**

Note: Sleep Quantity refers to the objectively measured average nightly sleep level (in min/night) for the week preceding the decision. Only compliant subjects are included, but variation in *Sleep Quantity* reflects different degrees of sleep restriction or rest in subjects. Response times (in milliseconds) are averaged over the set of trials (Easy vs Hard trials separated). Bayes scores are the total number of Bayesian correct responses out of 20 for the set of trials (Easy vs Hard trials separated). Two observations on each of the 119 subjects exist (1st and 2nd experiment session) and estimation include random effects error structure at the subject level.

*, **, *** indicate significance at the .10, .05, and .01 level, respectively, for the 2-tailed test.

Table A3 estimates the direct impact of self-reported sleepiness (Karolinska Sleepiness Scores) on average response time and Bayes scores in the set of Hard and Easy trials. As in the main manuscript, these results indicate that sleepiness (which is significantly increased in both the SD and MM treatment conditions) leads to longer response times. These estimates also indicate that increase sleepiness lowers the total number of Bayesian choices in both Hard and Easy trials. The main text found this to be true only for Hard trials. In the Easy trials, the results indicate either no impact of sleepiness on Bayes scores, or possibly an improvement in Bayes scores, which is consistent with the hypothesis that sleepier subjects are less likely to *incorrectly* use a deliberative thinking process (and therefore increase the chance of a Bayesian error) on the Easy trials.

Table A3: Self-Report Sleepiness impacts on Response Times and Bayes Scores

Coefficients (standard errors in parenthesis)

N=238 (119 subjects)	Response Time estimations		Bayes Scores Estimations	
	Dep Var Average RT (ms) Easy Trials	Dep Var Average RT (ms) Hard Trials	Dep Var Bayes Score Easy Trials	Dep Var Bayes Score Hard Trials
Constant	723.86 (45.43)***	823.72 (52.25)***	16.83 (.79)***	13.98 (.87)***
<i>2nd Decision Session</i>	-116.16 (22.16)***	-179.36 (25.67)***	1.83 (.39)***	.99 (.44)**
<i>Ksleepy</i>	12.31 (6.73)*	10.16 (7.77)	-.15 (.12)	-.26 (.13)**
Chi-squared	34.64***	54.69***	26.07***	10.82***

Note: *, **, *** indicate significance at the .10, .05, and .01 level, respectively, for the 2-tailed test.

Response times (in milliseconds) are averaged over the set of trials (Easy vs Hard trials separated). Bayes scores are the total number of Bayesian correct responses out of 20 for the set of trials (Easy vs Hard trials separated). Two observations on each of the 119 subjects exist (1st and 2nd experiment session) and estimation include random effects error structure at the subject level. Results are similar if including an additional control for the *2nd Decision Session*. Coefficient estimates on *2nd Decision Session* indicate lower average response times and increased Bayes Scores in both Hard and Easy trials for subjects who are being administered the tasks for the second time (the first administration would have been at the end of week 1 of the protocol, whereas the 2nd administration is at the end of week 3). *Ksleepy* still significantly increases average RT on Easy trials and lower Bayes scores in Hard trials. While qualitatively similar, the *Ksleepy* estimates on average RT on Hard trials and Bayes scores on Easy trials lose precision and are no longer statistically significant.

Table A4: Instrumental Variables Estimation

(data pooled across the 20 trials of Easy or Hard choices for a given subject for each Session)

Predictors of *Ksleepy* (self-report sleepiness)

1st Stage regression:

Constant	3.80 (.26)***
2 nd Decision Session	-.30 (.20)
Female	.87 (.22)***
<i>SD</i>	2.62 (.25)***
<i>MM</i>	1.19 (.35)***
<i>SD*MM</i>	-1.20 (.40)***
Morning Session	-.53 (.25)**
Morning* <i>MM</i>	-.05 (.41)
F(7,230)	30.23***

2nd Stage Estimations of Response Times and Bayes Scores

(uses predicted value of *Ksleepy* from 1st stage regression)

N=238 (119 subjects)	Response Time estimations		Bayes Scores Estimations	
	Dep Var Average RT (ms) Easy Trials	Dep Var Average RT (ms) Hard Trials	Dep Var Bayes Score Easy Trials	Dep Var Bayes Score Hard Trials
Constant	699.73 (54.21)***	791.78 (62.04)***	15.99 (1.17)***	15.26 (1.29)***
2 nd Decision Session	-114.02 (21.38)***	-176.52 (25.67)***	1.91 (.38)***	.88 (.46)*
<i>Ksleepy</i>	16.49 (8.83)*	15.70 (9.84)	-.0002 (.18)	-.49 (.22)**
Chi-squared	32.20***	55.56***	25.26***	12.99***

*, **, *** indicate significance at the .10, .05, and .01 level, respectively, for the 2-tailed test.

Robust standard errors are clustered on subject.