



# HHS Public Access

Author manuscript

*Adv Neural Inf Process Syst.* Author manuscript; available in PMC 2018 January 01.

Published in final edited form as:

*Adv Neural Inf Process Syst.* 2017 December ; 30: 5973–5981.

## Action Centered Contextual Bandits

**Kristjan Greenewald,**

Department of Statistics, Harvard University

**Ambuj Tewari,**

Department of Statistics, University of Michigan

**Predrag Klasnja, and**

School of Information, University of Michigan

**Susan Murphy**

Departments of Statistics and Computer Science, Harvard University

### Abstract

Contextual bandits have become popular as they offer a middle ground between very simple approaches based on multi-armed bandits and very complex approaches using the full power of reinforcement learning. They have demonstrated success in web applications and have a rich body of associated theoretical guarantees. Linear models are well understood theoretically and preferred by practitioners because they are not only easily interpretable but also simple to implement and debug. Furthermore, if the linear model is true, we get very strong performance guarantees. Unfortunately, in emerging applications in mobile health, the time-invariant linear model assumption is untenable. We provide an extension of the linear model for contextual bandits that has two parts: baseline reward and treatment effect. We allow the former to be complex but keep the latter simple. We argue that this model is plausible for mobile health applications. At the same time, it leads to algorithms with strong performance guarantees as in the linear model setting, while still allowing for complex nonlinear baseline modeling. Our theory is supported by experiments on data gathered in a recently concluded mobile health study.

## 1 Introduction

In the theory of sequential decision-making, contextual bandit problems (Tewari & Murphy, 2017) occupy a middle ground between multi-armed bandit problems (Bubeck & Cesa-Bianchi, 2012) and full-blown reinforcement learning (usually modeled using Markov decision processes along with discounted or average reward optimality criteria (Sutton & Barto, 1998; Puterman, 2005)). Unlike bandit algorithms, which cannot use any side-information or context, contextual bandit algorithms can learn to map the context into appropriate actions. However, contextual bandits do not consider the impact of actions on the evolution of future contexts. Nevertheless, in many practical domains where the impact of the learner’s action on future contexts is limited, contextual bandit algorithms have shown

great promise. Examples include web advertising (Abe & Nakamura, 1999) and news article selection on web portals (Li et al., 2010).

An influential thread within the contextual bandit literature models the expected reward for any action in a given context using a linear mapping from a  $d$ -dimensional context vector to a real-valued *reward*. Algorithms using this assumption include LinUCB and Thompson Sampling, for both of which regret bounds have been derived. These analyses often allow the context sequence to be chosen adversarially, but require the linear model, which links rewards to contexts, to be time-invariant. There has been little effort to extend these algorithms and analyses when the data follow an unknown nonlinear or time-varying model.

In this paper, we consider a particular type of non-stationarity and non-linearity that is motivated by problems arising in mobile health (mHealth). Mobile health is a fast developing field that uses mobile and wearable devices for health care delivery. These devices provide us with a real-time stream of dynamically evolving contextual information about the user (location, calendar, weather, physical activity, internet activity, etc.). Contextual bandit algorithms can learn to map this contextual information to a set of available intervention options (e.g., whether or not to send a medication reminder). However, human behavior is hard to model using stationary, linear models. We make a fundamental assumption in this paper that is quite plausible in the mHealth setting. In these settings, there is almost always a “do nothing” action usually called action 0. The expected reward for this action is the *baseline reward* and it can change in a very non-stationary, non-linear fashion. However, the *treatment effect* of a non-zero action, i.e., the incremental change over the baseline reward due to the action, can often be plausibly modeled using standard stationary, linear models.

We show, both theoretically and empirically, that the performance of an appropriately designed action-centered contextual bandit algorithm is agnostic to the high model complexity of the baseline reward. Instead, we get the same level of performance as expected in a stationary, linear model setting. Note that it might be tempting to make the entire model non-linear and non-stationary. However, the sample complexity of learning very general non-stationary, non-linear models is likely to be so high that they will not be useful in mHealth where data is often noisy, missing, or collected only over a few hundred decision points.

We connect our algorithm design and theoretical analysis to the real world of mHealth by using data from a pilot study of HeartSteps, an Android-based walking intervention. HeartSteps encourages walking by sending individuals contextually-tailored suggestions to be active. Such suggestions can be sent up to five times a day—in the morning, at lunchtime, mid-afternoon, at the end of the workday, and in the evening—and each suggestion is tailored to the user’s current context: location, time of day, day of the week, and weather. HeartSteps contains two types of suggestions: suggestions to go for a walk, and suggestions to simply move around in order to disrupt prolonged sitting. While the initial pilot study of HeartSteps micro-randomized the delivery of activity suggestions (Klasnja et al., 2015; Liao et al., 2015), delivery of activity suggestions is an excellent candidate for the use of contextual bandits, as the effect of delivering (vs. not) a suggestion at any given time is likely to be

strongly influenced by the user's current context, including location, time of day, and weather.

This paper's main contributions can be summarized as follows. We introduce a variant of the standard linear contextual bandit model that allows the baseline reward model to be quite complex while keeping the treatment effect model simple. We then introduce the idea of using action centering in contextual bandits as a way to decouple the estimation of the above two parts of the model. We show that action centering is effective in dealing with time-varying and non-linear behavior in our model, leading to regret bounds that scale as nicely as previous bounds for linear contextual bandits. Finally, we use data gathered in the recently conducted HeartSteps study to validate our model and theory.

### 1.1 Related Work

Contextual bandits have been the focus of considerable interest in recent years. Chu et al. (2011) and Agrawal & Goyal (2013) have examined UCB and Thompson sampling methods respectively for linear contextual bandits. Works such as Seldin et al. (2011), Dudik et al. (2011) considered contextual bandits with fixed policy classes. Methods for reducing the regret under complex reward functions include the nonparametric approach of May et al. (2012), the "contextual zooming" approach of Slivkins (2014), the kernel-based method of Valko et al. (2013), and the sparse method of Bastani & Bayati (2015). Each of these approaches has regret that scales with the complexity of the overall reward model including the baseline, and requires the reward function to remain constant over time.

## 2 Model and Problem Setting

Consider a contextual bandit with a baseline (zero) action and  $N$  non-baseline arms (actions or treatments). At each time  $t = 1, 2, \dots$ , a context vector  $\bar{s}_t \in \mathbb{R}^{d'}$  is observed, an action  $a_t \in \{0, \dots, N\}$  is chosen, and a reward  $r_t(a_t)$  is observed. The bandit learns a mapping from a state vector  $s_{t,a_t}$  depending on  $\bar{s}_t$  and  $a_t$  to the expected reward  $r_t(s_{t,a_t})$ . The state vector  $s_{t,a_t} \in \mathbb{R}^d$  is a function of  $a_t$  and  $\bar{s}_t$ . This form is used to achieve maximum generality, as it allows for infinite possible actions so long as the reward can be modeled using a  $d$ -dimensional  $s_{t,a}$ . In the most unstructured case with  $N$  actions, we can simply encode the reward with a  $d = Nd'$  dimensional  $s_{t,a_t}^T = [I(a_t=1)\bar{s}_t^T, \dots, I(a_t=N)\bar{s}_t^T]$  where  $I(\cdot)$  is the indicator function.

For maximum generality, we assume the context vectors are chosen by an adversary on the basis of the history  $\mathcal{H}_{t-1}$  of arms  $a_\tau$  played, states  $\bar{s}_\tau$ , and rewards  $r_\tau(\bar{s}_\tau, a_\tau)$  received up to time  $t-1$ , i.e.,

$$\mathcal{H}_{t-1} = \{a_\tau, \bar{s}_\tau, r_\tau(\bar{s}_\tau, a_\tau), i=1, \dots, N, \tau=1, \dots, t-1\}.$$

Consider the model  $E[r_t(\bar{s}_t, a_t) | \bar{s}_t, a_t] = \bar{f}_t(\bar{s}_t, a_t)$ , where  $\bar{f}_t$  can be decomposed into a fixed component dependent on action and a time-varying component that does not depend on action:

$$E[r_t(\bar{s}_t, a_t) | \bar{s}_t, a_t] = \bar{f}_t(\bar{s}_t, a_t) = f(s_{t,a_t})I(a_t > 0) + g_t(\bar{s}_t),$$

where  $\bar{f}_t(\bar{s}_t, 0) = g_t(\bar{s}_t)$  due to the indicator function  $I(a_t > 0)$ . Note that the optimal action depends in no way on  $g_t$ , which merely confounds the observation of regret. We hypothesize that the regret bounds for such a contextual bandit asymptotically depend only on the complexity of  $f$ , not of  $g_t$ . We emphasize that we do not require any assumptions about or bounds on the complexity or smoothness of  $g_t$ , allowing  $g_t$  to be arbitrarily nonlinear and to change abruptly in time. These conditions create a partially agnostic setting where we have a simple model for the interaction but the baseline cannot be modeled with a simple linear function. In what follows, for simplicity of notation we drop  $\bar{s}_t$  from the argument for  $r_t$ , writing  $r_t(a_t)$  with the dependence on  $\bar{s}_t$  understood.

In this paper, we consider the linear model for the reward difference at time  $t$ :

$$r_t(a_t) - r_t(0) = f(s_{t,a_t})I(a_t > 0) + n_t = s_{t,a_t}^T \theta I(a_t > 0) + n_t \quad (1)$$

where  $n_t$  is zero-mean sub-Gaussian noise with variance  $\sigma^2$  and  $\theta \in \mathbb{R}^d$  is a vector of coefficients. The goal of the contextual bandit is to estimate  $\theta$  at every time  $t$  and use the estimate to decide which actions to take under a series of observed contexts. As is common in the literature, we assume that both the baseline and interaction rewards are bounded by a constant for all  $t$ .

The task of the action-centered contextual bandit is to choose the probabilities  $\pi(a, t)$  of playing each arm  $a$ , at time  $t$  so as to maximize expected differential reward

$$\begin{aligned} [r_t(a_t) - r_t(0) | \mathcal{H}_{t-1}, s_{t,a}] &= \sum_{a=0}^N \pi(a, t) [r_t(a) - r_t(0) | \mathcal{H}_{t-1}, s_{t,a}] \\ &= \sum_{a=0}^N \pi(a, t) s_{t,a}^T \theta I(a > 0). \end{aligned} \quad (2)$$

This task is closely related to obtaining a good estimate of the reward function coefficients  $\theta$ .

## 2.1 Probability-constrained optimal policy

In the mHealth setting, a contextual bandit must choose at each time point whether to deliver to the user a behavior-change intervention, and if so, what type of intervention to deliver. Whether or not an intervention, such as an activity suggestion or a medication reminder, is sent is a critical aspect of the user experience. If a bandit sends too few interventions to a user, it risks the user's disengaging with the system, and if it sends too many, it risks the user's becoming overwhelmed or desensitized to the system's prompts. Furthermore, standard contextual bandits will eventually converge to a policy that maps most states to a near-100% chance of sending or not sending an intervention. Such regularity could not only worsen the user's experience, but ignores the fact that users have changing routines and

cannot be perfectly modeled. We are thus motivated to introduce a constraint on the size of the probabilities of delivering an intervention. We constrain  $0 < \pi_{\min} \leq 1 - \mathbb{P}(a_t = 0 | \bar{s}_t) \leq \pi_{\max} < 1$ , where  $\mathbb{P}(a_t = 0 | \bar{s}_t)$  is the conditional bandit-chosen probability of delivering an intervention at time  $t$ . The constants  $\pi_{\min}$  and  $\pi_{\max}$  are not learned by the algorithm, but chosen using domain science, and might vary for different components of the same mHealth system. We constrain  $\mathbb{P}(a_t = 0 | \bar{s}_t)$ , not each  $\mathbb{P}(a_t = i | \bar{s}_t)$ , as which intervention is delivered is less critical to the user experience than being prompted with an intervention in the first place. User habituation can be mitigated by implementing the nonzero actions ( $a = 1, \dots, N$ ) to correspond to several *types* or *categories* of messages, with the exact message sent being randomized from a set of differently worded messages.

Conceptually, we can view the bandit as pulling two arms at each time  $t$ : the probability of sending a message (constrained to lie in  $[\pi_{\min}, \pi_{\max}]$ ) and which message to send if one is sent. While these probability constraints are motivated by domain science, these constraints also enable our proposed action-centering algorithm to effectively orthogonalize the baseline and interaction term rewards, achieving sublinear regret in complex scenarios that often occur in mobile health and other applications and for which existing approaches have large regret.

Under this probability constraint, we can now derive the optimal policy with which to compare the bandit. The policy that maximizes the expected reward (2) will play the optimal action

$$a_t^* = \arg \max_{i \in \{0, \dots, N\}} s_{t,i}^T \theta I(i > 0),$$

with the highest allowed probability. The remainder of the probability is assigned as follows. If the optimal action is nonzero, the optimal policy will then play the zero action with the remaining probability (which is the minimum allowed probability of playing the zero action). If the optimal action is zero, the optimal policy will play the nonzero action with the highest expected reward

$$\bar{a}_t^* = \arg \max_{i \in \{1, \dots, N\}} s_{t,i}^T \theta$$

with the remaining probability, i.e.  $\pi_{\min}$ . To summarize, under the constraint

$1 - \pi_t^*(0, t) \in [\pi_{\min}, \pi_{\max}]$ , the expected reward maximizing policy plays arm  $a_t$  with probability  $\pi^*(a, t)$ , where

$$\begin{aligned} \text{If } a_t^* \neq 0: & \pi^*(a_t^*, t) = \pi_{\max}, \quad \pi^*(0, t) = 1 - \pi_{\max}, \quad \pi^*(a, t) = 0 \forall a \neq 0, a_t^* \\ \text{If } a_t^* = 0: & \pi^*(0, t) = 1 - \pi_{\min}, \quad \pi^*(\bar{a}_t^*, t) = \pi_{\min}, \quad \pi^*(a, t) = 0 \forall a \neq 0, \bar{a}_t^*. \end{aligned} \quad (3)$$

### 3 Action-centered contextual bandit

Since the observed reward always contains the sum of the baseline reward and the differential reward we are estimating, and the baseline reward is arbitrarily complex, the main challenge is to isolate the differential reward at each time step. We do this via an action-centering trick, which randomizes the action at each time step, allowing us to construct an estimator whose expectation is proportional to the differential reward  $r_t(a_t) - r_t(0)$ , where  $a_t$  is the nonzero action chosen by the bandit at time  $t$  to be randomized against the zero action. For simplicity of notation, we set the probability of the bandit taking nonzero action  $\mathbb{P}(a_t > 0)$  to be equal to  $1 - \pi_t(0, t) = \pi_t$ .

#### 3.1 Centering the actions - an unbiased $r_t(a_t) - r_t(0)$ estimate

To determine a policy, the bandit must learn the coefficients  $\theta$  of the model for the differential reward  $r_t(a_t) - r_t(0) = s_{t,a_t}^T \theta$  as a function of  $a_t$ . If the bandit had access at each time  $t$  to the differential reward  $r_t(a_t) - r_t(0)$ , we could estimate  $\theta$  using a penalized least-squares approach by minimizing

$$\arg \min_{\theta} \sum_{t=1}^T (r_t(a_t) - r_t(0) - \theta^T s_{t,a_t})^2 + \lambda \|\theta\|_2^2$$

over  $\theta$ , where  $r_t(a)$  is the reward under action  $a$  at time  $t$  (Agrawal & Goyal, 2013). This corresponds to the Bayesian estimator when the reward is Gaussian. Although we have only access to  $r_t(a_t)$ , not  $r_t(a_t) - r_t(0)$ , observe that given  $a_t$ , the bandit randomizes to  $a_t = a_t$  with probability  $\pi_t$  and  $a_t = 0$  otherwise. Thus

$$\begin{aligned} \mathbb{E}[(I(a_t > 0) - \pi_t)r_t(a_t) | \mathcal{H}_{t-1}, \bar{a}_t, \bar{s}_t] &= \pi_t(1 - \pi_t)r_t(\bar{a}) - (1 - \pi_t)\pi_t r_t(0) \\ &= \pi_t(1 - \pi_t)(r_t(\bar{a}_t) - r_t(0)). \end{aligned} \quad (4)$$

Thus  $(I(a_t > 0) - \pi_t)r_t(a_t)$ , which only uses the observed  $r_t(a_t)$ , is proportional to an unbiased estimator of  $r_t(a_t) - r_t(0)$ . Recalling that  $a_t$  and  $\pi_t$  are both known since they are chosen by the bandit at time  $t$ , we create the estimate of the differential reward between  $a_t$  and action 0 at time  $t$  as

$$\hat{r}_t(\bar{a}_t) = (I(a_t > 0) - \pi_t)r_t(a_t).$$

The corresponding penalized weighted least-squares estimator for  $\theta$  using  $\hat{r}_t(a_t)$  is the minimizer of

$$\begin{aligned} \sum_{t=1}^T \pi_t(1-\pi_t)(\hat{r}_t(\bar{a}_t)/(\pi_t(1-\pi_t))-\theta^T s_{t,\bar{a}_t})^2 + \|\theta\|_2^2 &= \sum_{t=1}^T \frac{(\hat{r}_t(\bar{a}_t))^2}{\pi_t(1-\pi_t)} - 2\hat{r}_t(\bar{a}_t)\theta^T s_{t,\bar{a}_t} + \pi_t(1-\pi_t)(\theta^T s_{t,\bar{a}_t})^2 + \|\theta\|_2^2 \\ &= c - 2\theta^T \hat{b} + \theta^T B\theta + \|\theta\|_2^2, \end{aligned}$$

(5)

where for simplicity of presentation we have used unit penalization  $\|\theta\|_2^2$  and

$$\hat{b} = \sum_{t=1}^T (I(a_t > 0) - \pi_t) s_{t,\bar{a}_t} r_t(a_t), \quad B = I + \sum_{t=1}^T \pi_t(1-\pi_t) s_{t,\bar{a}_t} s_{t,\bar{a}_t}^T.$$

The weighted least-squares weights are  $\pi_t(1-\pi_t)$ , since

$$\text{var} \left[ \frac{\hat{r}_t(\bar{a}_t)}{\pi_t(1-\pi_t)} \mid \mathcal{H}_{t-1}, \bar{a}_t, \bar{s}_t \right] = \frac{\text{var}[\hat{r}_t(\bar{a}_t) \mid \mathcal{H}_{t-1}, \bar{a}_t, \bar{s}_t]}{(\pi_t(1-\pi_t))^2}$$

and the standard deviation of  $\hat{r}_t(\bar{a}_t) = (I(a_t > 0) - \pi_t)r_t(a_t)$  given  $\mathcal{H}_{t-1}, \bar{a}_t, \bar{s}_t$  is of order  $g(\bar{s}_t) = \mathcal{O}(1)$ . The minimizer of (5) is  $\hat{\theta} = B^{-1}\hat{b}$ .

### 3.2 Action-Centered Thompson Sampling

As the Thompson sampling approach generates probabilities of taking an action, rather than selecting an action, Thompson sampling is particularly suited to our regression approach. We follow the basic framework of the contextual Thompson sampling approach presented by Agrawal & Goyal (2013), extending and modifying it to incorporate our action-centered estimator and probability constraints.

The critical step in Thompson sampling is randomizing the model coefficients according to the prior  $\mathcal{N}(\hat{\theta}, \nu^2 B^{-1})$  for  $\theta$  at time  $t$ . A  $\theta' \sim \mathcal{N}(\hat{\theta}, \nu^2 B^{-1})$  is generated, and the action  $a_t$  chosen to maximize  $s_{t,a}^T \theta'$ . The probability that this procedure selects any action  $a$  is determined by the distribution of  $\theta'$ ; however, it may select action 0 with a probability not in the required range  $[1 - \pi_{\max}, 1 - \pi_{\min}]$ . We thus introduce a two-step hierarchical procedure. After generating the random  $\theta'$ , we instead choose the *nonzero*  $a_t$  maximizing the expected reward

$$\bar{a}_t = \arg \max_{a \in \{1, \dots, N\}} s_{t,a}^T \theta'.$$

Then we randomly determine whether to take the nonzero action, choosing  $a_t = \bar{a}_t$  with probability

$$\pi_t = \mathbb{P}(a_t > 0) = \max(\pi_{\min}, \min(\pi_{\max}, \mathbb{P}(s_{t,\bar{a}_t}^T \tilde{\theta} > 0))), \quad (6)$$

and  $a_t = 0$  otherwise, where  $\tilde{\theta} \sim \mathcal{N}(\hat{\theta}, \nu^2 B^{-1})$ .  $\mathbb{P}(s_{t,\tilde{a}}^T \tilde{\theta} > 0)$  is the probability that the expected relative reward  $s_{t,\tilde{a}}^T \tilde{\theta}$  of action  $\tilde{a}_t$  is higher than zero for  $\tilde{\theta} \sim \mathcal{N}(\hat{\theta}, \nu^2 B^{-1})$ . This probability is easily computed using the normal CDF. Finally the bandit updates  $\hat{b}$ ,  $B$  and computes an updated  $\hat{\theta} = B^{-1} \hat{b}$ . Our action-centered Thompson sampling algorithm is summarized in Algorithm 1.

### Algorithm 1

#### Action-Centered Thompson Sampling

- 
- 1: Set  $B = I$ ,  $\hat{\theta} = 0$ ,  $\hat{b} = 0$ , choose  $[\pi_{\min}, \pi_{\max}]$ .
  - 2: **for**  $t = 1, 2, \dots$  **do**
  - 3:   Observe current context  $\bar{s}_t$  and form  $s_{t,a}$  for each  $a \in \{1, \dots, N\}$ .
  - 4:   Randomly generate  $\theta' \sim \mathcal{N}(\hat{\theta}, \nu^2 B^{-1})$ .
  - 5:   Let

$$\tilde{a}_t = \arg \max_{a \in \{1, \dots, N\}} s_{t,a}^T \theta'.$$

- 6:   Compute probability  $\pi_t$  of taking a nonzero action according to (6).
- 7:   Play action  $a_t = \tilde{a}_t$  with probability  $\pi_t$ , else play  $a_t = 0$ .
- 8:   Observe reward  $r_t(a_t)$  and update  $\hat{\theta}$

$$B = B + \pi_t(1 - \pi_t)s_{t,\tilde{a}_t}^T s_{t,\tilde{a}_t}, \quad \hat{b} = \hat{b} + s_{t,\tilde{a}_t}(I(a_t > 0) - \pi_t)r_t(a_t), \quad \hat{\theta} = B^{-1} \hat{b}.$$

**end for**

---

## 4 Regret analysis

Classically, the regret of a bandit is defined as the difference between the reward achieved by taking the optimal actions  $a_t^*$ , and the expected reward received by playing the arm  $a_t$  chosen by the bandit

$$\text{regret}_{\text{classical}}(t) = s_{t,a_t^*}^T \theta - s_{t,a_t}^T \theta, \quad (7)$$

where the expectation is taken conditionally on  $a_t, s_{t,a_t}^T, \mathcal{H}_{t-1}$ . For simplicity, let  $\pi_t^* = 1 - \pi_t^*(0, t)$  be the probability that the optimal policy takes a nonzero action, and recall that  $\pi_t = 1 - \pi_t(0, t)$  is the probability the bandit takes a nonzero action. The probability constraint implies that the optimal policy (3) plays the optimal arm with a probability bounded away from 0 and 1, hence definition (7) is no longer meaningful. We can instead create a regret that is the difference in expected rewards conditioned on  $a_t, \pi_t, s_{t,a_t}^T, \mathcal{H}_{t-1}$ , but not on the randomized action  $a_t$ :



$$\text{regret}(t) = \pi_t^* s_{t,\bar{a}_t}^T \theta - \pi_t s_{t,\bar{a}_t}^T \theta \quad (8)$$

where we have recalled that given  $\bar{p}$ , the bandit plays action  $a_t = \bar{a}_t$  with probability  $\pi_t$  and plays the  $a_t = 0$  with differential reward 0 otherwise. The action-centered contextual bandit attempts to minimize the cumulative regret  $\mathcal{R}(T) = \sum_{t=1}^T \text{regret}(t)$  over horizon  $T$ .

#### 4.1 Regret bound for Action-Centered Thompson Sampling

In the following theorem we show that with high probability, the probability-constrained Thompson sampler has low regret relative to the optimal probability-constrained policy.

**Theorem 1**—*Consider the action-centered contextual bandit problem, where  $\bar{f}_t$  is potentially time varying, and  $\bar{s}_t$  at time  $t$  given  $\mathcal{H}_{t-1}$  is chosen by an adversary. Under this regime, the total regret at time  $T$  for the action-centered Thompson sampling contextual bandit (Algorithm 1) satisfies*

$$\mathcal{R}(T) \leq C \left( \frac{d^2}{\varepsilon} \sqrt{T^{1+\varepsilon}} (\log(Td) \log \frac{1}{\delta}) \right)$$

with probability at least  $1 - 3\delta/2$ , for any  $0 < \varepsilon < 1$ ,  $0 < \delta < 1$ . The constant  $C$  is in the proof.

Observe that this regret bound does not depend on the number of actions  $N$ , is sublinear in  $T$ , and scales only with the complexity  $d$  of the interaction term, not the complexity of the baseline reward  $g$ . Furthermore,  $\varepsilon = 1/\log(T)$  can be chosen giving a regret of order  $O(d^2 \sqrt{T})$ .

This bound is of the same order as the baseline Thompson sampling contextual bandit in the adversarial setting when the baseline is identically zero (Agrawal & Goyal, 2013). When the baseline can be modeled with  $d'$  features where  $d' > d$ , our method achieves  $O(d^2 \sqrt{T})$  regret whereas the standard Thompson sampling approach has  $O((d+d')^2 \sqrt{T})$  regret. Furthermore, when the baseline reward is time-varying, the worst case regret of the standard Thompson sampling approach is  $O(T)$ , while the regret of our method remains  $O(d^2 \sqrt{T})$ .

#### 4.2 Proof of Theorem 1 - Decomposition of the regret

We will first bound the regret (8) at time  $t$ .

$$\text{regret}(t) = \pi_t^* s_{t,\bar{a}_t}^T \theta - \pi_t s_{t,\bar{a}_t}^T \theta = (\pi_t^* - \pi_t) (s_{t,\bar{a}_t}^T \theta) + \pi_t^* (s_{t,\bar{a}_t}^T \theta - s_{t,\bar{a}_t}^T \theta) \quad (9)$$

$$\leq (\pi_t^* - \pi_t)(s_{t,\bar{a}_t}^T \theta) + (s_{t,\bar{a}_t^*}^T \theta - s_{t,\bar{a}_t}^T \theta), \quad (10)$$

where the inequality holds since  $(s_{t,\bar{a}_t}^T \theta - s_{t,\bar{a}_t^*}^T \theta) \geq 0$  and  $0 < \pi_t^* < 1$  by definition. Then

$$\mathcal{R}(T) = \sum_{t=1}^T \text{regret}(t) \leq \underbrace{\sum_{t=1}^T (\pi_t^* - \pi_t)(s_{t,\bar{a}_t}^T \theta)}_I + \underbrace{\sum_{t=1}^T (s_{t,\bar{a}_t^*}^T \theta - s_{t,\bar{a}_t}^T \theta)}_{II}$$

Observe that we have decomposed the regret into a term  $I$  that depends on the choice of the randomization  $\pi_t$  between the zero and nonzero action, and a term  $II$  that depends only on the choice of the potential nonzero action  $\bar{a}_t$  prior to the randomization. We bound  $I$  using concentration inequalities, and bound  $II$  using arguments paralleling those for standard Thompson sampling.

**Lemma 1**—*Suppose that the conditions of Theorem 1 apply. Then with probability at least  $1 - \frac{\delta}{2}$ ,  $I \leq C \sqrt{d^3 T \log(Td) \log(1/\delta)}$  for some constant  $C$  given in the proof.*

**Lemma 2**—*Suppose that the conditions of Theorem 1 apply. Then term  $II$  can be bounded as*

$$II = \sum_{t=1}^T (s_{t,\bar{a}_t^*}^T \theta - s_{t,\bar{a}_t}^T \theta) \leq C' \left( \frac{d^2}{\varepsilon} \sqrt{T^{1+\varepsilon}} \log \frac{1}{\delta} \log(Td) \right)$$

where the inequality holds with probability at least  $1 - \delta$ .

The proofs are contained in Sections D and E in the supplement respectively. In the derivation, the “pseudo-actions”  $\bar{a}_t$  that Algorithm 1 chooses prior to the  $\pi_t$  baseline-nonzero randomization correspond to the actions in the standard contextual bandit setting. Note that  $I$  involves only  $\bar{a}_t$ , not  $\bar{a}_t^*$ , hence it is not surprising that the bound is smaller than that for  $II$ . Combining Lemmas 1 and 2 via the union bound gives Theorem 1.

## 5 Results

### 5.1 Simulated data

We first conduct experiments with simulated data, using  $N = 2$  possible nonzero actions. In each experiment, we choose a true reward generative model  $r_t(s, a)$  inspired by data from the HeartSteps study (for details see Section A.1 in the supplement), and generate two length  $T$  sequences of state vectors  $s_{t,a} \in \mathbb{R}^{NK}$  and  $\bar{s}_t \in \mathbb{R}^L$ , where the  $\bar{s}_t$  are iid Gaussian and  $s_{t,a}$  is formed by stacking columns  $\mathbb{I}(a = i)[1; \bar{s}_t]$  for  $i = 1, \dots, N$ . We consider both nonlinear and nonstationary baselines, while keeping the treatment effect models the same. The bandit

under evaluation iterates through the  $T$  time points, at each choosing an action and receiving a reward generated according to the chosen model. We set  $\pi_{\min} = 0.2$ ,  $\pi_{\max} = 0.8$ .

At each time step, the reward under the optimal policy is calculated and compared to the reward received by the bandit to form the regret  $\text{regret}(t)$ . We can then plot the cumulative regret

$$\text{cumulative regret}(t) = \sum_{\tau=1}^t \text{regret}(\tau).$$

In the first experiment, the baseline reward is nonlinear. Specifically, we generate rewards using  $r_t(s_{t,a_t}, \bar{s}_t, a_t) = \theta^T s_{t,a_t} + 2\mathcal{I}(|\bar{s}_t| < 0.8) + n_t$  where  $n_t = \mathcal{N}(0, 1)$  and  $\theta \in \mathbb{R}^8$  is a fixed vector listed in supplement section A.1. This simulates the quite likely scenario that for a given individual the baseline reward is higher for small absolute deviations from the mean of the first context feature, i.e. rewards are higher when the feature at the decision point is “near average”, with reward decreasing for abnormally high or low values. We run the benchmark Thompson sampling algorithm (Agrawal & Goyal, 2013) and our proposed action-centered Thompson sampling algorithm, computing the cumulative regrets and taking the median over 500 random trials. The results are shown in Figure 1, demonstrating linear growth of the benchmark Thompson sampling algorithm and significantly lower, sublinear regret for our proposed method.

We then consider a scenario with the baseline reward  $g_t(\cdot)$  function changing in time. We generate rewards as  $r_t(s_{t,a_t}, \bar{s}_t, a_t) = \theta^T s_{t,a_t} + \eta_t^T \bar{s}_t + n_t$  where  $n_t = \mathcal{N}(0, 1)$ ,  $\theta$  is a fixed vector as above, and  $\eta_t \in \mathbb{R}^7$ ,  $\bar{s}_t$  are generated as smoothly varying Gaussian processes (supplement Section A.1). The cumulative regret is shown in Figure 2, again demonstrating linear regret for the baseline approach and significantly lower sublinear regret for our proposed action-centering algorithm as expected.

## 5.2 HeartSteps study data

The HeartSteps study collected the sensor and weather-based features shown in Figure 3 at 5 decision points per day for each study participant. If the participant was available at a decision point, a message was sent with constant probability 0.6. The sent message could be one of several activity or anti-sedentary messages chosen by the system. The reward for that message was defined to be  $\log(0.5 + x)$  where  $x$  is the step count of the participant in the 30 minutes following the suggestion. As noted in the introduction, the baseline reward, i.e. the step count of a subject when no message is sent, does not only depend on the state in a complex way but is likely dependent on a large number of unobserved variables. Because of these unobserved variables, the mapping from the observed state to the reward is believed to be strongly time-varying. Both these characteristics (complex, time-varying baseline reward function) suggest the use of the action-centering approach.

We run our contextual bandit on the HeartSteps data, considering the binary action of whether or not to send a message at a given decision point based on the features listed in

Figure 3 in the supplement. Each user is considered independently, for maximum personalization and independence of results. As above we set  $\pi_{\min} = 0.2$ ,  $\pi_{\max} = 0.8$ .

We perform offline evaluation of the bandit using the method of Li et al. (2011). Li et al. (2011) uses the sequence of states, actions, and rewards in the data to form an near-unbiased estimate of the average expected reward achieved by each algorithm, averaging over all users. We used a total of 33797 time points to create the reward estimates. The resulting estimates for the improvement in average reward over the baseline randomization, averaged over 100 random seeds of the bandit algorithm, are shown in Figure 4 of the supplement with the proposed action-centering approach achieving the highest reward. Since the reward is logarithmic in the number of steps, the results imply that the benchmark Thompson sampling approach achieves an average 1.6% increase in step counts over the non-adaptive baseline, while our proposed method achieves an increase of 3.9%.

## 6 Conclusion

Motivated by emerging challenges in adaptive decision making in mobile health, in this paper we proposed the action-centered Thompson sampling contextual bandit, exploiting the randomness of the Thompson sampler and an action-centering approach to orthogonalize out the baseline reward. We proved that our approach enjoys low regret bounds that scale only with the complexity of the interaction term, allowing the baseline reward to be arbitrarily complex and time-varying.

## Acknowledgments

This work was supported in part by grants R01 AA023187, P50 DA039838, U54EB020404, R01 HL125440 NHLBI/NIA, NSF CAREER IIS-1452099, and a Sloan Research Fellowship.

## References

- Abbasi-Yadkori, Yasin, Pál, Dávid, Szepesvári, Csaba. Improved algorithms for linear stochastic bandits. *Advances in Neural Information Processing Systems*. 2011:2312–2320.
- Abe, Naoki, Nakamura, Atsuyoshi. Learning to optimally schedule internet banner advertisements. *Proceedings of the Sixteenth International Conference on Machine Learning*; Morgan Kaufmann Publishers Inc; 1999. p. 12-21.
- Agrawal, Shipra, Goyal, Navin. Thompson sampling for contextual bandits with linear payoffs. *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*; 2013. p. 127-135.
- Auer, Peter, Cesa-Bianchi, Nicolo, Freund, Yoav, Schapire, Robert E. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*. 2002; 32(1):48–77.
- Bastani, Hamsa, Bayati, Mohsen. Online decision-making with high-dimensional covariates. 2015 Available at SSRN 2661896.
- Bubeck, Sébastien, Cesa-Bianchi, Nicolo. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*. 2012; 5(1):1–122.
- Chu, Wei, Li, Lihong, Reyzin, Lev, Schapire, Robert E. Contextual bandits with linear payoff functions. *International Conference on Artificial Intelligence and Statistics*; 2011. p. 208-214.
- Dudik, Miroslav, Hsu, Daniel, Kale, Satyen, Karampatziakis, Nikos, Langford, John, Reyzin, Lev, Zhang, Tong. Efficient optimal learning for contextual bandits. *Proceedings of the Twenty-Seventh Conference Annual Conference on Uncertainty in Artificial Intelligence*; AUAI Press; 2011. p. 169-178.

- Klasnja, Predrag, Hekler, Eric B., Shiffman, Saul, Boruvka, Audrey, Almirall, Daniel, Tewari, Ambuj, Murphy, Susan A. Microrandomized trials: An experimental design for developing just-in-time adaptive interventions. *Health Psychology*. Dec; 2015 34(Suppl):1220–1228.
- Li, Lihong, Chu, Wei, Langford, John, Schapire, Robert E. A contextual-bandit approach to personalized news article recommendation. *Proceedings of the 19th International Conference on World Wide Web*; ACM; 2010. p. 661-670.
- Li, Lihong, Chu, Wei, Langford, John, Wang, Xuanhui. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. *Proceedings of the fourth ACM international conference on Web search and data mining*; ACM; 2011. p. 297-306.
- Liao, Peng, Klasnja, Predrag, Tewari, Ambuj, Murphy, Susan A. Sample size calculations for micro-randomized trials in mhealth. *Statistics in medicine*. 2015
- May, Benedict C., Korda, Nathan, Lee, Anthony, Leslie, David S. Optimistic Bayesian sampling in contextual-bandit problems. *The Journal of Machine Learning Research*. 2012; 13(1):2069–2106.
- Puterman, Martin L. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons; 2005.
- Seldin, Yevgeny, Auer, Peter, Shawe-Taylor, John S., Ortner, Ronald, Laviolette, François. PAC-Bayesian analysis of contextual bandits. *Advances in Neural Information Processing Systems*. 2011:1683–1691.
- Slivkins, Aleksandrs. Contextual bandits with similarity information. *The Journal of Machine Learning Research*. 2014; 15(1):2533–2568.
- Sutton, Richard S., Barto, Andrew G. *Reinforcement learning: An introduction*. MIT Press; 1998.
- Tewari, Ambuj, Murphy, Susan A. From ads to interventions: Contextual bandits in mobile health. In: Reh, Jim, Murphy, Susan A., Kumar, Santosh, editors. *Mobile Health: Sensors, Analytic Methods, and Applications*. Springer; 2017.
- Valko, Michal, Korda, Nathan, Munos, Rémi, Flaounas, Ilias, Cristianini, Nello. Finite-time analysis of kernelised contextual bandits. *Uncertainty in Artificial Intelligence*. 2013:654.

## A HeartSteps feature list

Figure 3 shows the features available to the bandit in the HeartSteps study dataset, and Figure 4 shows the estimated average regret results with errorbars.

### A.1 Simulation model

Figure 5 shows the coefficients  $\theta$  used in the main text simulations. The coefficients shown in the figure associated with the first action are obtained via a linear regression analysis of the binary action (sending or not sending a message) HeartSteps intervention data, and the coefficients for the second action are a simple modification of those.

For the time varying simulation, Gaussian processes were used to generate the reward coefficient sequence  $\eta_t$  and the state sequence  $\bar{s}_t$ . We used Gaussian processes since if  $\eta_t$  is IID, then the baseline reward becomes an IID random variable, making the baseline reward not time varying.

We used the Gaussian process

$$\eta_t = \sqrt{1 - \rho^2} \eta_{t-1} + \rho \epsilon_t$$

where  $\eta_0 = 17$ ,  $n_t \sim \mathcal{N}(0, I_7)$ , and  $\rho = 0.1$ . The state sequence  $\bar{s}_t$  was generated in the same manner.

## B Definitions

In order to proceed with the proof of Theorem 1, we make the following definitions.

### Definition 1

Define a filtration  $\mathcal{F}_{t-1} = \{\mathcal{H}_{t-1}, \bar{s}_t\}$  as the union of the history and current context.

### Definition 2

Let

$$z_{t,a} = \sqrt{s_{t,a}^T B(t)^{-1} s_{t,a}},$$

for all  $a = 1, \dots, N$ .

### Definition 3

Define  $\ell(T) = R \sqrt{d \log(T^3) \log(1/\delta)} + 1$ ,  $v = R \sqrt{\frac{24}{\epsilon} d \log(1/\delta)}$ , and  $g(T) = \sqrt{4d \log(Td)v + \ell(T)}$ .

We divide the arms  $> 0$  into saturated and unsaturated actions.

### Definition 4 (Saturated vs. unsaturated actions)

Any arm  $> 0$  for which  $g(T)z_{t,\bar{a}} < \ell(T)z_{t,\bar{a}^*}$  is called a saturated arm. If an arm is not saturated, it is called unsaturated. Let  $\mathcal{C}(t) \subseteq \{1, \dots, N\}$  be the subset of saturated arms at time  $t$ .

Observe that the optimal arm  $^*$  is unsaturated by definition.

We can now state the required concentration events and present bounds on the probability they occur.

## B.1 Concentration events

### Definition 5

Let  $E^{\mu}(t)$  be the event that for all  $= 1, \dots, N$

$$|s_{t,\bar{a}}^T \hat{\theta}_t - s_{t,\bar{a}}^T \theta| \leq \ell(T)z_{t,\bar{a}}.$$

Similarly, let  $E^{\theta}(t)$  be the event that for all  $= 1, \dots, N$

$$|s_{t,\bar{a}}^T \theta'_t - s_{t,\bar{a}}^T \hat{\theta}_t| \leq \sqrt{4d \log(Td) v z_{t,\bar{a}}}$$

and  $E_0^\theta(t)$  be the corresponding event that for all  $i = 1, \dots, N$

$$|s_{t,\bar{a}}^T \tilde{\theta}_t - s_{t,\bar{a}}^T \hat{\theta}_t| \leq \sqrt{4d \log(Td) v z_{t,\bar{a}}}$$

We can bound the probabilities of the events  $E^\theta(t)$ ,  $E^\theta(t)_0$ , and  $E^\mu(t)$  in the following lemmas. Observe that by definition  $\mathbb{P}(E^\theta(t) | \mathcal{F}_{t-1}) = \mathbb{P}(E^\theta(t)_0 | \mathcal{F}_{t-1})$ .

### Lemma 3 (Agrawal & Goyal (2013))

For all  $t$ , and possible filtrations  $\mathcal{F}_{t-1}$ ,  $\mathbb{P}(E^\theta(t) | \mathcal{F}_{t-1}) \geq 1 - \frac{1}{T^2}$ .

For  $E^\mu(t)$  we have

### Lemma 4

For all  $t$ ,  $0 < \delta < 1$ ,  $\mathbb{P}(E^\mu(t)) \geq 1 - \frac{\delta}{T^2}$ .

The proof is given in Section G.

## B.2 Supermartingales

### Definition 6 (Supermartingale)

A sequence of random variables  $(Y_t; t \geq 0)$  is called a supermartingale corresponding to a filtration  $\mathcal{F}_t$  if, for all  $t$ ,  $Y_t$  is  $\mathcal{F}_t$ -measurable, and

$$\mathbb{E}[Y_t - Y_{t-1} | \mathcal{F}_{t-1}] \leq 0$$

for all  $t \geq 1$ .

### Lemma 5 (Azuma-Hoeffding inequality)

If for all  $t = 1, \dots, T$  a supermartingale  $(Y_t; t \geq 0)$  corresponding to filtration  $\mathcal{F}_t$  satisfies  $|Y_t - Y_{t-1}| \leq c_t$  for some constants  $c_t$ , then for any  $a > 0$

$$\mathbb{P}(Y_T - Y_0 \geq a) \leq e^{-\frac{a^2}{2 \sum_{t=1}^T c_t^2}}.$$

## C Preliminary results

### C.1 Lemma 7: Probability of choosing a saturated action $\bar{a}_t \in C(t)$

#### Lemma 6 (Agrawal & Goyal (2013) Lemma 2)

For any filtration  $\mathcal{F}_{t-1}$  such that  $E^u(t)$  is true,

$$\mathbb{P}(s_{t,\bar{a}_t}^T \theta' > s_{t,\bar{a}_t}^T \theta + \ell(T) z_{t,\bar{a}_t} | \mathcal{F}_{t-1}) \geq \frac{1}{4e \sqrt{\pi T \varepsilon}}.$$

We can now prove the following.

#### Lemma 7

For any filtration  $\mathcal{F}_{t-1}$  such that  $E^u(t)$  is true,

$$\mathbb{P}(\bar{a}_t \in C(t) | \mathcal{F}_{t-1}) \leq \frac{1}{p} \mathbb{P}(\bar{a}_t \notin C(t) | \mathcal{F}_{t-1}) + \frac{1}{pT^2},$$

where  $p = \frac{1}{4e \sqrt{\pi T \varepsilon}}$ .

**Proof**—Recall that  $\bar{a}_t$  is the action with the largest value of  $s_{t,\bar{a}_t}^T \theta'$ . Hence, if  $s_{t,\bar{a}_t}^T \theta'$  is larger than  $s_{t,i}^T \theta'$  for all  $i \in C(t)$ , then  $\bar{a}_t$  is one of the unsaturated actions. Hence

$$\mathbb{P}(\bar{a}_t \notin C(t) | \mathcal{F}_{t-1}) \geq \mathbb{P}(s_{t,\bar{a}_t}^T \theta' > s_{t,i}^T \theta', \forall i \in C(t) | \mathcal{F}_{t-1}). \quad (11)$$

We know that by definition all saturated arms  $i \in C(t)$  have  $g(T) z_{t,j} < \ell(T) z_{t,\bar{a}_t}$ . Given an  $\mathcal{F}_{t-1}$  such that  $E^u(t)$  holds, we have that either  $E^\theta(t)$  is false or for all  $i \in C(t)$

$$s_{t,i}^T \theta' \leq s_{t,i}^T \theta + g(T) z_{t,i} \leq s_{t,\bar{a}_t}^T \theta + \ell(T) z_{t,\bar{a}_t}$$

implying

$$\begin{aligned} & \mathbb{P}(s_{t,\bar{a}_t}^T \theta' > s_{t,i}^T \theta', \forall j \in C(t) | \mathcal{F}_{t-1}) \\ & \geq \mathbb{P}(s_{t,\bar{a}_t}^T \theta' > s_{t,\bar{a}_t}^T \theta + \ell(T) z_{t,\bar{a}_t} | \mathcal{F}_{t-1}) - \mathbb{P}(E^\theta(t) | \mathcal{F}_{t-1}) \\ & \geq p - \frac{1}{T^2}. \end{aligned}$$



where we have used the definitions of  $E^u(t)$ ,  $E^\theta(t)$ , and the last inequality follows from Lemma 6 and Lemma 4. Substituting into (11) gives

$$\mathbb{P}(\bar{a}_t \notin C(t) | \mathcal{F}_{t-1}) + \frac{1}{T^2} \geq p,$$

and

$$\frac{\mathbb{P}(\bar{a}_t \in C(t) | \mathcal{F}_{t-1})}{\mathbb{P}(\bar{a}_t \notin C(t) | \mathcal{F}_{t-1}) + \frac{1}{T^2}} \leq \frac{1}{p}.$$

## C.2 Lemma 9 - Bound on $\sum_t z_t$

### Lemma 8

For  $z_{t,a} = \sqrt{s_{t,a}^T B(t)^{-1} s_{t,a}}$ , we have that

$$\sum_{t=1}^T z_{t,\bar{a}_t} \leq \frac{5}{C_\pi} \sqrt{dT \log T},$$

where  $C_\pi = \sqrt{\min(\pi_{\min}(1-\pi_{\max}), \pi_{\max}(1-\pi_{\min}))}$  is a constant.

**Proof**—We apply the following lemma from Auer et al. (2002) and Chu et al. (2011).

### Lemma 9

Let  $A_t = I + \sum_{t=1}^T x_t x_t^T$ , where  $x_t \in \mathbb{R}^d$  is a sequence of vectors. Then, defining

$\sigma_t = \sqrt{x_t^T A_t^{-1} x_t}$ , we have

$$\sum_{t=1}^T \sigma_t \leq 5 \sqrt{dT \log T}.$$

To apply this to  $\sum_t z_t$ , let  $x_t = \sqrt{\pi_t(1-\pi_t)} s_{t,\bar{a}_t}$ . Then

$A_t = I + \sum_{t=1}^T (\pi_t(1-\pi_t)) s_{t,\bar{a}_t} s_{t,\bar{a}_t}^T = B_t$ , and we have

$$\sigma_t = \sqrt{x_t^T A_t^{-1} x_t} = \sqrt{\pi_t(1-\pi_t)} \sqrt{s_{t,\bar{a}_t}^T B_t^{-1} s_{t,\bar{a}_t}} = \sqrt{\pi_t(1-\pi_t)} z_{t,\bar{a}_t}.$$

Applying Lemma 9 we thus have

$$\sum_{t=1}^T z_{t,\bar{a}_t} \leq \max_t \left( \frac{1}{\sqrt{\pi_t(1-\pi_t)}} \right) \sum_{t=1}^T \sigma_t \leq \frac{5}{C_\pi} \sqrt{dT \log T},$$

where  $C_\pi = \sqrt{\min(\pi_{\min}(1-\pi_{\max}), \pi_{\max}(1-\pi_{\min}))}$  is a constant.

## D Proof of Lemma 1 - term I

**Proof**—We know that by definition of the optimal policy,  $(\pi_t^* - \pi_t) s_{t,\bar{a}_t}^T \theta \geq 0$ . Hence under event  $E^{\mathcal{U}}(t)$ ,

$$\begin{aligned} (\pi_t^* - \pi_t) s_{t,\bar{a}_t}^T \theta &\leq \mathbb{P} \left( \text{sign}(s_{t,\bar{a}_t}^T \theta') \neq \text{sign}(s_{t,\bar{a}_t}^T \theta) \right) |s_{t,\bar{a}_t}^T \theta| \\ &\leq \min [ |s_{t,\bar{a}_t}^T \theta|, \mathbb{P}(\text{sign}(s_{t,\bar{a}_t}^T \theta') \neq \text{sign}(s_{t,\bar{a}_t}^T \theta)) ] \\ &\leq (\ell(T) + \sqrt{4d \log(Td)v}) z_{t,\bar{a}_t} + 1 - \mathbb{P}(E_0^\theta(t)). \end{aligned}$$

Substituting in the definitions of  $\ell(T)$ ,  $v$  and the bound in Lemma 3 on  $\mathbb{P}(E_0^\theta(t))$ , we have

$$\begin{aligned} (\pi_t^* - \pi_t) s_{t,\bar{a}_t}^T \theta &\leq \left( R \sqrt{d \log(T^3) \log(1/\delta)} + 1 + \sqrt{4d \log(Td)R} \sqrt{\frac{24}{\varepsilon} d \log(1/\delta)} \right) z_{t,\bar{a}_t} + \frac{1}{T^2} \\ &\leq C \sqrt{\frac{d^2}{\varepsilon} \log(1/\delta)} z_{t,\bar{a}_t} + \frac{1}{T^2}. \end{aligned}$$

Summing over  $t$  and recalling that by Lemma 9  $\sum_{t=1}^T z_{t,\bar{a}_t} \leq \frac{5}{C_\pi} \sqrt{dT \log T}$ , we have that under event  $E^{\mathcal{U}}(t)$

$$\begin{aligned} I &= \sum_{t=1}^T (\pi_t^* - \pi_t) s_{t,\bar{a}_t}^T \theta \\ &\leq \frac{C}{C_\pi} \sqrt{d^3 T \log(Td) \log(1/\delta)}. \end{aligned}$$

Since the probability that  $E^{\mathcal{U}}(t)$  holds is at least  $1 - \frac{\delta}{T^2}$  by Lemma 4, the lemma results.

## E Proof of Lemma 2: Bound on term II

Before commencing the proof, we first state the following result from Abbasi-Yadkori et al. (2011).

**Lemma 10 (Abbasi-Yadkori et al. (2011))**

Let  $(\mathcal{F}'_t; t \geq 0)$  be a filtration,  $(m_t; t \geq 1)$  be an  $\mathbb{R}^d$ -valued stochastic process such that  $m_t$  is  $(\mathcal{F}'_{t-1})$ -measurable,  $(\eta_t; t \geq 1)$  be a real-valued martingale difference process such that  $\eta_t$  is  $(\mathcal{F}'_t)$ -measurable. For  $t \geq 0$ , define  $\xi_t = \sum_{\tau=1}^t m_\tau \eta_\tau$  and  $M_t = I_d + \sum_{\tau=1}^t m_\tau m_\tau^T$ , where  $I_d$  is the  $d$ -dimensional identity matrix. Assume  $\eta_t$  is conditionally  $R$ -sub-Gaussian.

Then, for any  $\delta' > 0$ ,  $t \geq 0$ , with probability at least  $1 - \delta'$ ,

$$\|\xi_t\|_{M_t^{-1}} \leq R \sqrt{d \log \left( \frac{t+1}{\delta'} \right)},$$

where  $\|\xi_t\|_{M_t^{-1}} = \sqrt{\xi_t^T M_t^{-1} \xi_t}$ .

We now prove Lemma 2.

**Proof**—Defining  $\text{regret}'(t) = (s_{t, \bar{a}_t^*}^T \theta - s_{t, \bar{a}_t}^T \theta) I(E^\mu(t))$ , we have the following lemma, which we prove in Section F.

**Lemma 11**

Let, for  $p = \frac{1}{4e \sqrt{\pi T^\varepsilon}}$ ,

$$X_t = \text{regret}'(t) - \frac{g(T)}{p} I(a(t) \notin C(t)) z_{t, \bar{a}_t^*} \quad (12)$$

$$Y_t = \sum_{w=1}^t X_w. \quad (13)$$

Then  $(Y_t; t = 0, \dots, T)$  is a super-martingale process with respect to filtration  $\mathcal{F}_t$ .

Given our results in Section G.1 and our concentration bounds, the proof is closely related to Agrawal & Goyal (2013) and is listed in Section F.

Using the definition of  $X_t$ , we have that

$$|Y_t - Y_{t-1}| \leq |X_t| \leq 1 + \frac{g(T)}{p} + \frac{2g(T)^2}{\ell(T)} + \frac{2g(T)}{pT^2} \leq \frac{8}{p} \frac{g(T)^2}{\ell(T)}. \text{ This allows us to apply the Azuma-Hoeffding inequality listed in Section B.2, giving that}$$

$$\begin{aligned} \sum_{t=1}^T \text{regret}'(t) &\leq \sum_{t=1}^T \left( \frac{g(T)}{p} I(\bar{a}_t \notin C(t)) z_{t, \bar{a}_t^*} \right) + \frac{2g(T)}{pT} + \frac{2g(T)^2}{\ell(T)} \sum_{t=1}^T z_{t, \bar{a}_t} + \frac{8}{p} \frac{g(T)^2}{\ell(T)} \sqrt{2T \log \frac{2}{\delta}} \\ &\leq \sum_{t=1}^T \left( \frac{g(T)^2}{\ell(T)} \frac{1}{p} I(\bar{a}_t \notin C(t)) z_{t, \bar{a}_t} \right) + \frac{2g(T)}{pT} + \frac{2g(T)^2}{\ell(T)} \sum_{t=1}^T z_{t, \bar{a}_t} + \frac{8}{p} \frac{g(T)^2}{\ell(T)} \sqrt{2T \log \frac{2}{\delta}} \\ &\leq \frac{g(T)^2}{\ell(T)} \frac{3}{p} \sum_{t=1}^T z_{t, \bar{a}_t} + \frac{2g(T)}{pT} + \frac{8}{p} \frac{g(T)^2}{\ell(T)} \sqrt{2T \log \frac{2}{\delta}}. \end{aligned}$$

with probability at least  $1 - \delta/2$ , where we recall that if  $t \notin C(t)$ , then  $g(T) z_{t, \bar{a}_t} \geq \ell(T) z_{t, \bar{a}_t^*}$ .

Substituting in the bound  $\sum_{t=1}^T z_{t, \bar{a}_t} \leq \frac{5}{c_\pi} \sqrt{dT \log T}$  from Lemma 9 and the definitions of  $g(T)$ ,  $p$ ,  $\ell(T)$ , we obtain that

$$\sum_{t=1}^T \text{regret}'(t) \leq \frac{C'}{C_\pi} \left( \frac{d^2}{\varepsilon} \sqrt{T^{1+\varepsilon}} \log \frac{1}{\delta} \log(Td) \right)$$

with probability at least  $1 - \frac{\delta}{2}$ , where  $C'$  is a constant. Recall that by Lemma 4,  $E^u(t)$  holds

for all  $t$  with probability at least  $1 - \delta/2$ , and that  $\text{regret}'(t) = (s_{t, \bar{a}_t^*}^T \theta - s_{t, \bar{a}_t}^T \theta)$  whenever  $E^u(t)$  holds. By the union bound we then have that

$$II = \sum_{t=1}^T (s_{t, \bar{a}_t^*}^T \theta - s_{t, \bar{a}_t}^T \theta) \leq \frac{C'}{C_\pi} \left( \frac{d^2}{\varepsilon} \sqrt{T^{1+\varepsilon}} \log \frac{1}{\delta} \log(Td) \right)$$

with probability at least  $1 - \delta$ . The lemma results.

## F Proof of Lemma 11

**Proof**—To prove that  $Y_t$  is a super-martingale by the definition above, we need to prove that for all  $1 \leq t \leq T$  and any  $\mathcal{F}_{t-1}$ ,  $\mathbb{E}[Y_t - Y_{t-1} | \mathcal{F}_{t-1}] \leq 0$ .

We first consider filtrations  $\mathcal{F}_{t-1}$  for which  $E^u(t)$  holds. By the definition of  $\rho$

$s_{t, \bar{a}_t}^T \theta' \geq s_{t, a_t^{**}}^T \theta'$ . Under  $E^\theta(t)$  and  $E^u(t)$  we then must have that for all  $i = 1, \dots, N$

$$\begin{aligned} s_{t, i}^T \theta &\geq s_{t, i}^T \theta' - g(T) z_{t, i} \\ &\geq s_{t, a_t^{**}}^T \theta' - g(T) z_{t, i} \\ &\geq s_{t, a_t^{**}}^T \theta - g(T) z_{t, \bar{a}_t^*} - g(T) z_{t, i}. \end{aligned}$$

Hence  $s_{t, \bar{a}_t^*}^T \theta - s_{t, \bar{a}_t}^T \theta \leq g(T)(z_{t, \bar{a}_t} + z_{t, \bar{a}_t^*})$ .

For  $\mathcal{F}_{t-1}$  such that  $E^\mu(t)$  holds, we then can write

$$\begin{aligned} [\text{regret}'(t)|\mathcal{F}_{t-1}] &= [(s_{t,\bar{a}_t^*}^T \theta - s_{t,\bar{a}_t}^T \theta)|\mathcal{F}_{t-1}] \\ &\geq [g(T)(z_{t,\bar{a}_t} + z_{t,\bar{a}_t^*})|\mathcal{F}_{t-1}] + \mathbb{P}(\overline{E^\theta}(t)) \\ &= g(T)z_{t,\bar{a}_t^*} \mathbb{P}(\bar{a}_t \in C(t)|\mathcal{F}_{t-1}) + g(T) \left[ \left( \frac{g(T)}{\ell(T)} z_{t,\bar{a}_t} I(\bar{a}_t \notin C(t)) \right) |\mathcal{F}_{t-1} \right] + g(T)[z_{t,\bar{a}_t}|\mathcal{F}_{t-1}] + \frac{1}{T^2}. \end{aligned}$$

where we have used the facts that  $\text{regret}'(t) \geq 1$ , the definition of unsaturated arms, and Lemma 4. Applying Lemma 7 and noting that since  $\min \text{eig}(B(t)) \geq 1$ ,  $z_{t,i} \leq \|s_{t,i}\|_2 \leq 1$ , we can show that

$$[\text{regret}'(t)|\mathcal{F}_{t-1}] \leq \frac{g(T)}{p} \mathbb{P}(\bar{a}_t \notin C(t)|\mathcal{F}_{t-1}) z_{t,\bar{a}_t^*} + \frac{2g(T)}{pT^2}. \quad (14)$$

By definition,  $\text{regret}'(t) = (s_{t,\bar{a}_t^*}^T \theta - s_{t,\bar{a}_t}^T \theta) I(E^\mu(t))$  is zero and the above inequality holds whenever  $E^\mu(t)$  is not true. Since we have considered both cases, the lemma is proved.

## G Proof of Lemma 4

**Proof**—We can apply Lemma 10 with  $m_t = \sqrt{\pi_t(1-\pi_t)} s_{t,\bar{a}_t}$ ,

$$\eta_t = \frac{\hat{r}_t(\bar{a}_t)}{\sqrt{\pi_t(1-\pi_t)}} - \sqrt{\pi_t(1-\pi_t)} s_{t,\bar{a}_t}^T \theta,$$

and with the filtration  $\mathcal{F}'_t = (\bar{s}_{\tau+1}, m_{\tau+1}, \eta_\tau : \tau \leq t)$  effectively containing all the available information up to the current time.  $\mathcal{F}'_{t-1}$  is measurable by definition, and in Section G.1 we show

### Lemma 12

Suppose that  $\eta_t$  is  $\mathbb{R}$  sub-Gaussian. Then  $\eta_t$  is a  $\mathcal{F}'_t$ -measurable,  $\mathbb{R}'$ -sub-Gaussian, martingale difference process where  $R' = \frac{R+2}{\sqrt{\pi_{\min}(1-\pi_{\max})}} + \sqrt{\pi_{\max}(1-\pi_{\min})}$ .

We then have

$$\begin{aligned}
M_t &= I_d + \sum_{\tau=1}^t m_\tau m_\tau^T = I_d + \sum_{\tau=1}^t \pi_\tau (1 - \pi_\tau) s_{\tau, \bar{a}_\tau} s_{\tau, \bar{a}_\tau}^T, \\
\xi_t &= \sum_{\tau=1}^t m_\tau \eta_\tau = \sum_{\tau=1}^t s_{\tau, \bar{a}_\tau} \left( \hat{r}_t(\bar{a}_t) - \pi_t (1 - \pi_t) s_{t, \bar{a}_t}^T \theta \right).
\end{aligned}$$

Observe that these are the two primary components of the contextual bandit, specifically,  $B_t = M_{t-1}$  and  $b_t - \mathbb{E}[b_t] = \xi_t$ . Hence,  $\hat{\theta}_t - \theta = M_{t-1}^{-1} (\xi_{t-1} - \theta)$ . Letting  $\|y\|_A = \sqrt{y^T A y}$  for any vector  $y$  and matrix  $A \in \mathbb{R}^{d \times d}$ , for all  $\delta > 0$  we have that since  $M_t$  is positive definite,

$$\begin{aligned}
|s_{\bar{a}, t}^T \hat{\theta} - s_{\bar{a}, t}^T \theta| &= |s_{\bar{a}, t}^T M_{t-1}^{-1} (\xi_{t-1} - \theta)| \\
&\leq \|s_{\bar{a}, t}\|_{M_{t-1}^{-1}} \|\xi_{t-1} - \theta\|_{M_{t-1}^{-1}} \\
&= \|s_{\bar{a}, t}\|_{B_t^{-1}} \|\xi_{t-1} - \theta\|_{B_t^{-1}}.
\end{aligned}$$

Applying Lemma 10, we have that for any  $\delta' > 0$ ,  $t \geq 1$ ,

$$\|\xi_{t-1}\|_{M_{t-1}^{-1}} \leq R' \sqrt{d \log \frac{t}{\delta'}}.$$

Then  $\|\xi_{t-1} - \theta\|_{M_{t-1}^{-1}} \leq R' \sqrt{d \log \frac{t}{\delta'}} + \|\theta\|_{M_{t-1}^{-1}} \leq R' \sqrt{d \log \frac{t}{\delta'}} + 1$ . Setting  $\delta' = \delta/T^2$  implies that with probability  $1 - \delta/T^2$ , for all  $t$ ,

$$|s_{\bar{a}, t}^T \hat{\theta} - s_{\bar{a}, t}^T \theta| \leq \|s_{\bar{a}, t}\|_{B_t^{-1}} \left( R' \sqrt{d \log(T^3) \log \frac{1}{\delta}} + 1 \right) = \ell(T) z_{t, \bar{a}}.$$

## G.1 Proof of Lemma 12: Martingale analysis of $\eta_t$

**Proof**—Recall

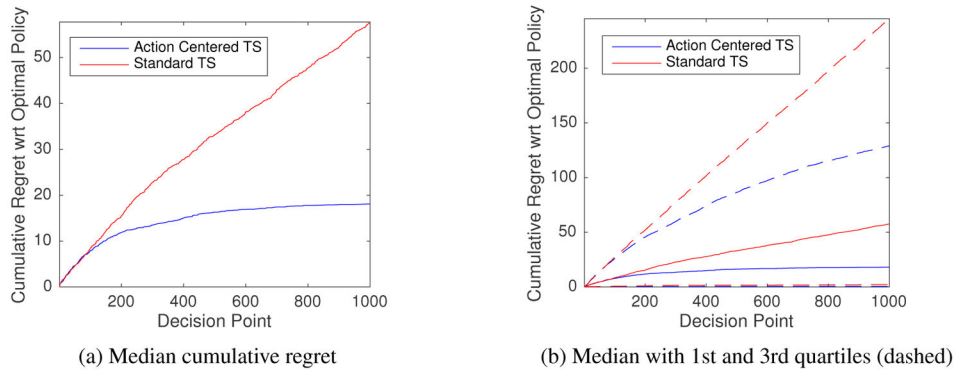
$$\begin{aligned}
|\eta_t| &= \left| \frac{\hat{r}_t(\bar{a}_t)}{\sqrt{\pi_t(1-\pi_t)}} - \sqrt{\pi_t(1-\pi_t)} s_{t,\bar{a}_t}^T \theta \right| \\
&= \left| \frac{(I(a_t > 0) - \pi_t) r_t(a_t)}{\sqrt{\pi_t(1-\pi_t)}} - \sqrt{\pi_t(1-\pi_t)} s_{t,\bar{a}_t}^T \theta \right| \\
&= \left| \frac{(I(a_t > 0) - \pi_t) (s_{t,a_t}^T \theta I(a_t > 0) + n_t + \bar{f}_t(\bar{s}_t))}{\sqrt{\pi_t(1-\pi_t)}} - \sqrt{\pi_t(1-\pi_t)} s_{t,\bar{a}_t}^T \theta \right| \\
&\leq \sqrt{\pi_t(1-\pi_t)} + \left| \frac{2+n_t}{\sqrt{\pi_t(1-\pi_t)}} \right|.
\end{aligned}$$

since the rewards are all bounded by one and the  $\pi_{\min} \leq \pi_t \leq \pi_{\max}$  are bounded. We have assumed that  $n_t$  is  $R$  sub-Gaussian. Since a bounded random variable  $|X| < b$  is  $b$  sub-Gaussian and the sum of independent  $b_1$  and  $b_2$  sub-Gaussian random variables is  $b_1 + b_2$  sub-Gaussian, we have that  $\eta_t$  is  $R' = \frac{R+2}{\sqrt{\pi_{\max}(1-\pi_{\min})}} + \sqrt{\pi_{\min}(1-\pi_{\max})}$  conditionally sub-Gaussian. Since  $\pi_{\min}, \pi_{\max}$  are bounded away from 0 and 1 by constants,  $R'$  is a constant.

Additionally, for all  $t$

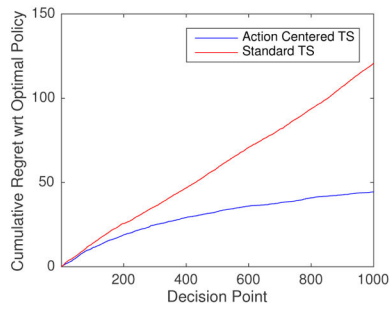
$$\begin{aligned}
[\eta_t | \mathcal{H}_{t-1}, \bar{a}_t, \bar{s}_t] &= \frac{[\hat{r}_t(\bar{a}_t) | \mathcal{H}_{t-1}, \bar{a}_t, \bar{s}_t]}{\sqrt{\pi_t(1-\pi_t)}} - \sqrt{\pi_t(1-\pi_t)} s_{t,\bar{a}_t}^T \theta \\
&= \frac{[(I(a_t > 0) - \pi_t) r_t(a_t) | \mathcal{H}_{t-1}, \bar{a}_t, \bar{s}_t]}{\sqrt{\pi_t(1-\pi_t)}} - \sqrt{\pi_t(1-\pi_t)} s_{t,\bar{a}_t}^T \theta \\
&= \frac{\pi_t(1-\pi_t) s_{t,\bar{a}_t}^T \theta}{\sqrt{\pi_t(1-\pi_t)}} - \sqrt{\pi_t(1-\pi_t)} s_{t,\bar{a}_t}^T \theta \\
&= 0,
\end{aligned}$$

where the third equality follows from (4). Thus  $E[\eta_t | \mathcal{H}_{t-1}, \bar{s}_t] = 0$  and  $\eta_t$  is a martingale difference process.




**Figure 1.** Nonlinear baseline reward  $g$ , in scenario with 2 nonzero actions and reward function based on collected HeartSteps data. Cumulative regret shown for proposed Action-Centered approach, compared to baseline contextual bandit, median computed over 100 random trials.





(a) Median cumulative regret



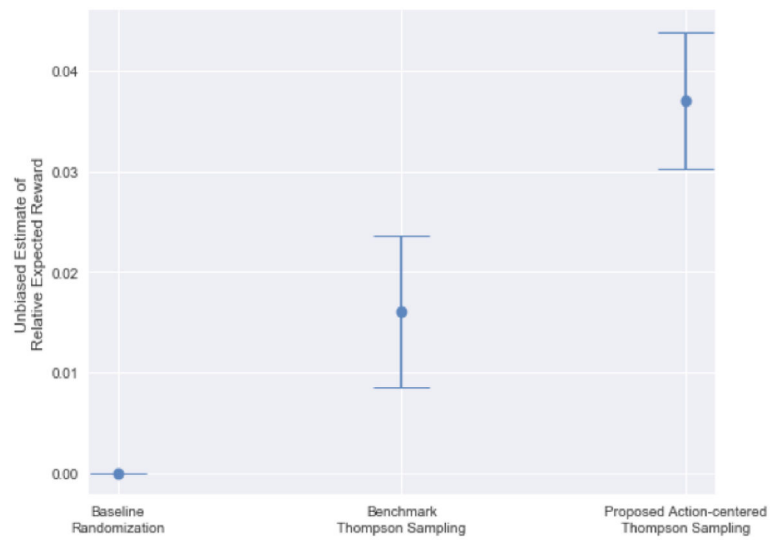
(b) Median with 1st and 3rd quartiles (dashed)

**Figure 2.** Nonstationary baseline reward  $g$ , in scenario with 2 nonzero actions and reward function based on collected HeartSteps data. Cumulative regret shown for proposed Action-Centered approach, compared to baseline contextual bandit, median computed over 100 random trials.

Feature	Description	Purpose	Interaction	Baseline Model
Number of messages sent	Total number of messages sent to user in prior week	Modeling habituation to intervention	Y	Y
Location indicator 1	1 if not at home or work, 0 o.w.	Location relevant to availability to walk	Y	Y
Location indicator 2	1 if at work, 0 o.w.		Y	Y
Step count variability	Historical standard deviation of step counts in 60 minute window surrounding decision point, taken over prior 7 days	Responsiveness in different times of day	Y	Y
Steps in prior 30 minutes	Step count in 30 minutes prior to decision point	Measure of recent activity		Y
Square root of steps yesterday	Square root of the total step count yesterday	Recent commitment/ engagement		Y
Outdoor Temperature	Degrees Celsius	Cold weather potentially less appealing		Y

**Figure 3.**

List of features available to the bandit in the HeartSteps experiment. The features available to model the action interaction (effect of sending an anti-sedentary message) and to model the baseline (reward under no action) are denoted via a “Y” in the corresponding column.



**Figure 4.**

Unbiased estimates of the average reward received by the benchmark Thompson sampling contextual bandit and the proposed action-centered Thompson sampling contextual bandit, relative to the reward received under the pre-specified HeartSteps randomization policy. Also shown are one standard deviation error bars for the computed estimates. The superior performance of the action-centering approach is indicative of its robustness to the high complexity of the baseline subject behavior.

Feature	Action 1 coef.	Action 2 coef.
Number of messages sent	.116	.116
Location indicator 1	-.275	.275
Location indicator 2	-.233	-.233
Step count variability	.0425	.0425

**Figure 5.**  
Effect coefficients, based on HeartSteps data, used for simulation reward model.