

Online Control For Adaptive Tapering Of Medications

Paula Gradu and Benjamin Recht

Abstract—We investigate adaptive protocols for the elimination or reduction of the use of medications or addictive substances. We formalize this problem as online optimization, minimizing the cumulative dose subject to constraints on individual well-being. We adapt a model of addiction from the psychology literature and show how it can be described by a class of linear time-invariant systems. For such systems, the optimal policy amounts to taking the smallest dose that maintains well-being. We derive a simple protocol based on integral control that requires no system identification, only needing approximate knowledge of the instantaneous dose response. This protocol is robust to model misspecification and is able to maintain an individual’s well-being during the tapering process. Numerical experiments demonstrate that the adaptive protocol outperforms non-adaptive methods in terms of both maintenance of well-being and rate of dose reduction.

I. INTRODUCTION

Tapering medications and assisting cessation of addictive substances are related challenges in health care. Relapse rates are high, and many return to drug use within weeks of entering treatment. Though no general protocols are equally valuable to all medications or all people, a personalized approach to cessation may allow caregivers flexibility to meet the diverse needs of a diverse care-seeking population. This paper investigates adaptive tapering protocols that enable individuals to self-regulate their cessation rate. In particular, we formalize tapering as an optimization problem of minimizing the cumulative dose subject to constraints on well-being.

Such a formulation requires modeling an individual’s dose-response dynamics, and we propose a mathematical formulation of the *opponent process* model of addiction due to Solomon [1]. As we discuss in Section II, an opponent process consists of two competing systemic reactions to treatment. The first system governs the “positive” effect with a rapid onset and fast subsequent decay. The second system governs “negative effects” with a long latency and slow decay. In Section III, we show that modeling opponent processes as LTI systems captures the qualitative behaviors specified in the psychology literature and allows us to analyze tapering as an optimal control problem.

In Section IV, we first derive a greedy policy that solves the optimal control problem. The dosage should be decreased by as much as possible while maintaining constraints on well-being, but no look-ahead or planning is needed to compute the dose. We provide a robust controller that maintains this greedy behavior and is optimal under various models of possible exogenous disturbances.

P. Gradu and B. Recht are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, USA pgradu@berkeley.edu, brecht@berkeley.edu.

In reality, the opponent process model is an approximation, and the particular model of an individual is unknown. A possible approach informed by control theoretic practice might involve learning a dynamics model of the individual, but such learning procedures require wide varying of inputs and would not be feasible for patients attempting drug cessation. Instead of leaning on system identification, we investigate the potential of “model-free” controllers for tapering in Section V. We analyze the potential of integral control, where the error signal is the deviation from a minimally acceptable level of well-being. Using methods from the online learning literature, we show that integral control robustly reduces the dosage while minimally violating constraints on well-being. Under further mild assumptions about the time between doses, we show that the integral controller monotonically reduces the dose to zero in finite time without violating the specified constraints.

In the numerical experiments of Section VI, we demonstrate that our adaptive protocol outperforms non-adaptive methods in terms of both maintenance of well-being and rate of dose reduction.

II. RELATED WORK

A. Tapering Protocols

Our work attempts to synthesize adaptive protocols that can be applied to both “deprescribing” medications via gradual tapering and to managing the symptoms of withdrawal from addictive substances. In both of these applications, most studies of tapering protocols have focused on non-personalized procedures. Mujika and Padilla [2] study tapering protocols for SSRI medications. They find that short tapers (2-4 weeks) have minimal benefits over abrupt discontinuation while longer tapers (multiple months) are more successful at minimizing withdrawal. More specifically they suggest a slow, hyperbolic reduction in dosage. Horowitz et al. [3], [4] propose a similar method (very slow, hyperbolic) for tapering of antipsychotic medication. This work is based on case studies and explicitly notes that there is no standard guideline for tapering antipsychotic medications.

Several studies have also investigated protocols for managing withdrawal from addictive substances. In a meta-analysis of opioid tapering protocols, Berna et al. [5] found that longer tapers are typically better. The survey by Fenton et al. [6] argues that most tapering protocols which focus primarily on getting through the acute withdrawal phase may be too rapid and have negative mental health consequences. In an observational study, Agnoli et al. [7] find that tapering of opioids is significantly associated with increased risk of overdose and mental health crisis. Henry et al. [8] attempt

to characterize patients' subjective experiences with opioid tapering in an effort to minimize negative tapering reactions.

B. Models of Drug Response and Tolerance

In the present work, we will build upon a popular model of addiction, proposed by Solomon [1], called the Opponent-Process theory of acquired motivation. In this model, the drug response is the result of an initial 'positive' effect with short lag and fast decay (A process) followed by a counter 'negative' effect with high latency and slow decay (B process). This formalizes the empirical qualitative observation that the withdrawal symptoms of a drug are characterized as opposite to its acute effects [9]. The shape of such a response is plotted in Figure 1. Koob and Moal [10] and Koob [11] extend the opponent process framework to account for a chronic deviation of the regulatory system from baseline under repeated administration of the drug, which is referred to as allostasis and is illustrated in Figure 2. In this work, we propose a simple mathematical formulation of the opponent process that captures the salient aspects of this model.




Fig. 1. Shape of an Instantaneous Opponent Process response as described by Solomon [1].




Fig. 2. Allostasis development due to overlay of opponent processes over time as described by Koob and Moal [10].

III. PROBLEM SETTING AND PRELIMINARIES

To mathematically formalize the opponent process model, we will denote an individual's well-being at any time as y_t . This formulation asserts that y_t is a concrete numerical quantity, but our methods can operate with imprecise quantification of well-being. For example, for modeling opiate withdrawal, y_t could be a value on the Clinical Opiate Withdrawal Scale [12] or the Subjective/Objective Opiate Withdrawal

Scales [13]. For modeling an individual's response to anti-depressants, y_t could model the number of Discontinuation-Emergent Signs and Symptoms (DESS) [14], a depression assessment like PHQ-9 [15], or a combination of both.

Let u_t denote the amount of a drug ingested at time t . We assume that y_t evolves according to an LTI system

$$y_{t+1} = \sum_{k=0}^t g(k)u_{t-k} + y_{t+1}^{\text{nat}} \quad (1)$$

where g is the impulse response. We focus our attention on linear models as they are sufficiently general to capture much of the qualitative behavior observed in the psychology literature, but we discuss potential extensions to nonlinear models in the Appendix. y^{nat} comprises the remaining dynamics that cannot be attributed to the drug under study. We will only lightly model this exogenous disturbance by asserting lower bounds on its values or bounds on its rate of change. The time $t = 0$ will be the moment an individual adopts a tapering protocol. The effect of all inputs before $t = 0$ and any measurement noise are implicitly represented in the "natural progression" signal, y^{nat} .

We formalize the goal of tapering as minimizing the total consumption of a drug subject to keeping well-being above some specified threshold. The total consumed drug is $\sum_{t=0}^{T-1} u_t$ and the constraint on well-being is $y_t \geq y_{\min}$ for all t from 1 to T .

A. Modeling Opponent Processes

We now turn to our formal model of an opponent process. Solomon [1] asserts that an opponent process reaction to a single dose is first an A-process where $y > 0$ and then a following B-process where $y < 0$. The following definition summarizes this property.

Definition 1 (Opponent Process): We say that g is an opponent process if there exists a time τ_0 such that $g(\tau) > 0$ when $\tau < \tau_0$ and $g(\tau) \leq 0$ if $\tau \geq \tau_0$.

In this work, we focus on a subclass of opponent processes that both induce some notion of tolerance or addiction and are sufficiently well-behaved to be "treatable."

Definition 2: An opponent process g is a *Linearly Progressing Opponent Process (LPOP)* if there additionally exists an $\alpha \in (0, 1)$ such that

$$g(t+1) \leq \alpha \cdot g(t) \quad \text{for } t < \tau_0 - 1$$

$$|g(t+1)| \geq \alpha \cdot |g(t)| \quad \text{for } t \geq \tau_0.$$

This definition quantifies opponent processes where the A-process decays exponentially, and the B-process does not decay as quickly as the A-process. For example, this would model a drug where the positive effects decay at some rate α , then become negative, trending towards some peak of withdrawal, and finally decay back to 0, potentially at a prolonged rate.

B. Examples of Opponent Processes

Consider a linear system with impulse response

$$g(t) = \sum_{\lambda} c_{\lambda} \lambda^t$$

where all of the $\lambda \in [0, 1)$ and the c_λ are real valued scalars. The terms where c_λ are positive correspond to effects that increase well-being. The terms where c_λ are negative decrease well-being. With this in mind, define $\Lambda_+ = \{\lambda : c_\lambda \geq 0\}$ and $\Lambda_- = \{\lambda : c_\lambda < 0\}$. If

- 1) $\max_{\lambda \in \Lambda_+} \lambda \leq \min_{\lambda' \in \Lambda_-} \lambda'$
- 2) $\sum_{\lambda} c_\lambda > 0$
- 3) There exists some τ_0 s.t. $g(\tau_0) \leq 0$

then the system is an LPOP.

Intuitively, this holds because the part of the system corresponding to positive effects decay more rapidly than all of the terms corresponding to negative effects. In particular, the second condition implies an initial positive response, and third condition implies that after a single dose, the well-being will eventually be negative.

The formal argument of why such systems are LPOPs proceeds by bounding $g(t+1)$ in terms of $g(t)$. Let $\Lambda_{++} := \max_{\lambda \in \Lambda_+} \lambda$ and $\Lambda_{--} := \min_{\lambda' \in \Lambda_-} \lambda'$. Then we have

$$\begin{aligned} g(t+1) &= \sum_{\lambda} c_\lambda \lambda^{t+1} \\ &= \sum_{\lambda_+ \in \Lambda_+} c_{\lambda_+} \lambda_+^{t+1} + \sum_{\lambda_- \in \Lambda_-} c_{\lambda_-} \lambda_-^{t+1} \\ &\leq \Lambda_{++} \cdot \sum_{\lambda_+ \in \Lambda_+} c_{\lambda_+} \lambda_+^t + \Lambda_{--} \cdot \sum_{\lambda_- \in \Lambda_-} c_{\lambda_-} \lambda_-^t \\ &\leq \Lambda_{--} \left(\sum_{\lambda_+ \in \Lambda_+} c_{\lambda_+} \lambda_+^t + \sum_{\lambda_- \in \Lambda_-} c_{\lambda_-} \lambda_-^t \right) \\ &= \Lambda_{--} g(t) \end{aligned}$$

Setting $\alpha \doteq \Lambda_{--}$, this implies that for $t < \tau_0 - 1$, $g(t+1) \leq \alpha g(t)$ and for $t \geq \tau_0$, $|g(t+1)| \geq \alpha |g(t)|$.

IV. OPTIMAL TAPERING PROTOCOL WITH SYSTEM KNOWLEDGE

In this section, we derive the optimal tapering protocol when the quantities described in the previous section are known to the care provider. In particular, we will show that when the system dynamics are known, the optimal control problem is equivalent to setpoint matching with setpoint y_{\min} .

We begin with the formal definition of a *maximally greedy* dose, the smallest dose that does not lead to an immediate violation of the constraint on y . The main theorem of this section states that a maximally greedy protocol results in the lowest cumulative dose.

Definition 3 (Maximally Greedy Dose): We say that a dose u_t is *maximally greedy* at time t if it either leads to $y_{t+1} = y_{\min}$ or is equal to 0.

Theorem 1: For any LPOP g , taking the maximally greedy dose at every time maintains $y_t \geq y_{\min}$ for all t with minimal cumulative dose.

Proof: Consider a dosing sequence u_t^* that maintains $y_t^* \geq y_{\min}$ for all t with minimal cumulative dose. We will prove the theorem by contradiction, assuming the sequence u_t^* is not maximally greedy. Then there must exist t_0 and

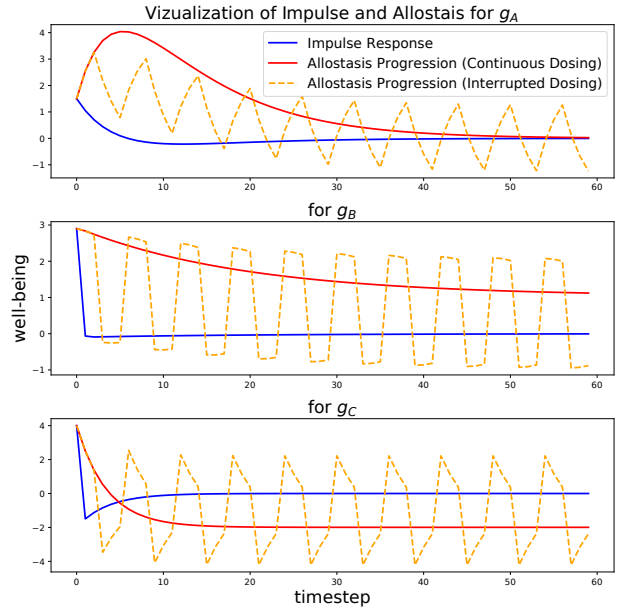


Fig. 3. Impulse responses and allostasis progression over 60 timesteps for three opponent processes, each with different time constants and relative effect magnitudes. The first two panels illustrate an opponent processes with immediate benefits and slowly accumulating tolerance, characteristic of therapeutic medications. The third panel captures an opponent process with immediate transient effect and also a rapid onset of a low value of well-being, more characteristic of substances of abuse.

$\epsilon > 0$ such that replacing $u_{t_0}^*$ with $u_{t_0}^* - \epsilon$ would result in a new $y'_{t_0+1} < y_{t_0+1}^*$ with $y'_{t_0+1} \geq y_{\min}$. Consider the modified dosing schedule

$$u'_t = \begin{cases} u_{t_0}^* - \epsilon & t = t_0 \\ u_{t_0+1}^* + \alpha \epsilon & t = t_0 + 1 \\ u_t^* & \text{otherwise} \end{cases}$$

and let y'_t be the resulting sequence. Since $u'_t = u_t^*$ for all $t < t_0$, we have that $y'_t = y_t^* \geq y_{\min}$ up to t_0 . Plugging in the effect of the dose modifications at t_0 and $t_0 + 1$ into the dynamics equation (1), we get:

$$\begin{aligned} y'_{t_0+1} &= y_{t_0+1}^* - g(t-t_0)\epsilon + g(t-t_0-1)\alpha\epsilon \\ &\geq y_{t_0+1}^* \\ &\geq y_{\min} \end{aligned} \quad (g \text{ LPOP})$$

since this holds for all $t > t_0$, we have that y'_t satisfies the constraints for all $1 \leq t \leq T$. But we note that u'_t has a smaller cumulative dose than u_t^* because

$$\sum_{t=1}^T u'_t = \sum_{t=1}^T u_t^* - \epsilon + \alpha\epsilon < \sum_{t=1}^T u_t^*.$$

This is a contradiction, completing the proof. ■

This theorem shows that long-term planning is unnecessary to optimally taper a medication. It always suffices to take the smallest dose necessary to maintain well being at the current time. The question remains if such myopic strategies remain possible when properties of the dynamics and dose responses are unknown.

Let us first assume that we know the impulse response g . Our first protocol considers the case when we have a lower bound on the values of y_t^{nat} .

Proposition 1: For any LPOP g , assume we know a lower bound $y_t^{\text{nat,lb}}$ such that $y_t^{\text{nat,lb}} \leq y_t^{\text{nat}}$ for all t . Then for any instance of y_t^{nat} that is greater than the prescribed lower bound, the dosing sequence

$$u_t = \max \left\{ 0, \frac{y_{\min} - y_{t+1}^{\text{nat,lb}} - \sum_{k=1}^t g(k)u_{t-k}}{g(0)} \right\} \quad (2)$$

maintains $y_t \geq y_{\min}$ for all t with minimal cumulative dose.

Proof: To prove this proposition, we first derive the maximally greedy dose at time t . This dose will require access to y_{t+1}^{nat} – we will replace this term with its lower bound and show the resulting protocol is optimal.

The maximally greedy dose at time t (given past doses) is

$$u_t^{\text{MG}} = \max \left\{ 0, \frac{y_{\min} - y_{t+1}^{\text{nat}} - \sum_{k=1}^t g(k)u_{t-k}}{g(0)} \right\} \quad (3)$$

To see why, observe that the dose is either 0 or satisfies

$$\begin{aligned} y_{t+1}^{\text{MG}} &= u_t^{\text{MG}} g(0) + \sum_{k=1}^t g(k)u_{t-k} + y_{t+1}^{\text{nat}} \\ &= y_{\min} - y_{t+1}^{\text{nat}} - \sum_{k=1}^t g(k)u_{t-k} + \sum_{k=1}^t g(k)u_{t-k} + y_{t+1}^{\text{nat}} \\ &= y_{\min} \end{aligned}$$

For any $y_{t+1}^{\text{nat}} \geq y_{t+1}^{\text{nat,lb}}$, the dose given by (2) is larger than the maximally greedy dose (i.e. $u_t \geq u_t^{\text{MG}}$) and therefore maintains $y_{t+1} \geq y_{t+1}^{\text{MG}} \geq y_{\min}$. This holds independently for all t (irrespective of past doses) and hence (2) maintains $y_t \geq y_{\min}$ for all t for any sequence y_t^{nat} such that $y_t^{\text{nat}} \geq y_t^{\text{nat,lb}}$ for all t .

Finally, since the sequence of inputs given in (3) is maximally greedy with respect to the lower bound sequence $y_y^{\text{nat,lb}}$, Theorem 1 implies that any dosing schedule with a strictly lower cumulative dose would have to violate one of the constraints. This means that u_t given by (2) results in the minimal cumulative dose that maintains feasibility across all admissible y_t^{nat} sequences (i.e. lower bounded by $y_t^{\text{nat,lb}}$). ■

There are two natural choices to lower bound y_t^{nat} . First, we may assume that $y_{t+1}^{\text{nat}} \geq y_t^{\text{nat}}$ for all t . In this case, the best lower bound candidate is $y_{t+1}^{\text{nat,lb}} = y_t^{\text{nat}}$ which we can compute via

$$y_t^{\text{nat}} = y_t - \sum_{k=0}^{t-1} g(k)u_{t-k-1}. \quad (4)$$

Second, we could bound the variation of y_t^{nat} , assuming there $L_{\text{nat}} \geq 0$ such that $y_{t+1}^{\text{nat}} \geq y_t^{\text{nat}} - L_{\text{nat}}$ for all t . In this case, the best lower bound is $y_{t+1}^{\text{nat,lb}} = y_t^{\text{nat}} - L_{\text{nat}}$, where we again compute y_t^{nat} via Eq. 4.

Note that the proof of Theorem 1 is valid even if the constraints on y or the impulse response g are time-varying. Hence, even under more general settings, setpoint matching

is the optimal approach to tapering. The doses for the time-varying setting are the same as in Proposition 1 but with $y_{\min}^{(t)}$ and/or g_t plugged into the formula. In de-medication, time-varying constraints are relevant when the patient can withstand more severe withdrawal effects at particular times of their lives or when they can no longer tolerate an initially agreed-upon baseline.

V. TAPERING BY INTEGRAL CONTROL

Computing the maximally greedy dose in the previous section required knowledge of the model of the underlying dynamics, g . In practice, we would like to avoid estimating g and instead derive an update rule requiring minimal system knowledge. Since we know that the optimal control law is setpoint matching and that integral action can match setpoints asymptotically, we analyze the performance of a simple integral controller on the tapering problem in this section. We demonstrate that integral control with appropriately chosen gains can reduce a dose while rarely violating the constraint on y_t . One need only know an approximate value of the immediate effect of a single dose, which is equal to the value $g(0)$. Moreover, if $g(t) \leq 0$ for all $t \geq 1$, then integral control results in a dose sequence monotonically decreases to zero in finite time while ensuring $y_t \geq y_{\min}$ for all t .

To proceed, let $(x)_+ := \max(x, 0)$ and $(x)_- := \min(x, 0)$. Consider the integral control law

$$u_t = \max \left\{ 0, u_{t-1} - K_+(y_t - y_{\min})_+ - K_-(y_t - y_{\min})_- \right\}. \quad (5)$$

This protocol is integral control if $K_+ = K_-$. However, it also allows for a different gain when y_t is above and below the desired set point y_{\min} . Using two gains allows for more conservative (i.e., higher) doses. The control action is clipped at 0 as it is impossible to take a negative dose.

The main theorem of this section bounds the *long-term constraint violation* [16], [17] of the integral controller. We aim to study how the average value of y_t compares to the constraint y_{\min} over time. The bound we derive ensures the overall sequence of states satisfies the constraint more strictly as time increases. This ensures that the running average of the observations up until a time T is greater than y_{\min} minus a penalty decaying at rate $\frac{1}{T}$.

Theorem 2: For any underlying y^{nat} , (5) ensures that for any T

$$\frac{\sum_{t=1}^T y_t}{T} \geq y_{\min} - \frac{y_0 - y_{\min}}{T},$$

provided $K_+ \leq g(0)^{-1}$ and $K_- \geq g(0)^{-1}$.

The proof of this theorem is broadly inspired from analyses of the finite time behavior of average constraint violation in regret minimization by Mannor and Tsitsiklis [17] and Mahdavi et al. [16]. To prove the theorem, we need the following technical lemma that lower bounds the instance-wise deviation of the y_t obtained from the dosing scheme (5) from y_{\min} if the gains K_+/K_- upper/lower bound $g(0)^{-1}$ respectively. The proof of the lemma is given in the Appendix.

Lemma 1: Adopt the convention that $u_{-1} = 0$. For any gains satisfying $K_+ \leq g(0)^{-1}$ and $K_- \geq g(0)^{-1}$, setting u_t according to (5) for $t \geq 0$ yields y_{t+1} satisfying

$$y_{t+1} \geq y_{\min} + (y_{t+1}^{\text{nat}} - y_t^{\text{nat}}) - \sum_{k=1}^t g(k)(u_{t-k-1} - u_{t-k}).$$

With this lemma, we can now prove Theorem 2.

Proof: [of Theorem 2] Using Lemma 1, we have that

$$\begin{aligned} \sum_{t=0}^T y_{t+1} &\geq (T+1)y_{\min} \\ &+ \sum_{t=0}^T ((y_{t+1}^{\text{nat}} - y_t^{\text{nat}}) - \sum_{k=1}^t g(k)(u_{t-k-1} - u_{t-k})) \\ &= (T+1)y_{\min} - y_0^{\text{nat}} + y_{T+1}^{\text{nat}} + \sum_{t=0}^T u_t g(T-t) \\ &= (T+1)y_{\min} + y_{T+1} - y_0 \end{aligned}$$

where the first equality follows by telescoping and the second by noting that $y_0^{\text{nat}} = y_0$. By cancelling y_{T+1} on both sides, we get

$$\sum_{t=0}^{T-1} y_{t+1} \geq T y_{\min} - (y_0 - y_{\min})$$

which can be immediately rearranged into the stated bound:

$$\frac{\sum_{t=1}^T y_t}{T} = \frac{\sum_{t=0}^{T-1} y_{t+1}}{T} \geq y_{\min} - \frac{y_0 - y_{\min}}{T}$$

Note that if it is undesirable to fluctuate around y_{\min} , we can run a padded integral controller adding δ to y_{\min} . Theorem 2 then implies

$$\frac{\sum_{t=1}^T y_t}{T} \geq y_{\min} + \delta - \frac{y_0 - y_{\min} - \delta}{T}.$$

Also note that, just as in the previous section, we can extend the integral controller (5) to time-varying constraints by plugging in $y_{\min}^{(t)}$. In this case, the guarantee from Theorem 2 would instead bound deviation from the average y_{\min} , i.e. from $T^{-1} \sum_{t=1}^T y_{\min}^{(t)}$.

We close this section by deriving stronger guarantees for a special class of LPOPs. Consider an opponent process with $\tau_0 = 1$. Requiring $\tau_0 = 1$ can be interpreted as an assumption on the discretization of the dose updating scheme (e.g., every day vs. every week). For example, we can turn a system g with $\tau_0 > 1$ into a system g' with $\tau_0' = 1$ by taking $g'(t) = \tau_0^{-1} \cdot \sum_{t'=t-\tau_0}^{(t+1)\cdot\tau_0-1} g(t')$.

Proposition 2: For any g with $\tau_0 = 1$, any initial u_0 which ensures $y_1 \geq y_{\min}$, and any natural progression sequence satisfying $y_{t+1}^{\text{nat}} \geq y_t^{\text{nat}} - g(t)u_0 + \delta/t$ for some (arbitrarily small) $\delta > 0$, the doses prescribed by (5): (i) are non-increasing, (ii) maintain $y_t \geq y_{\min}$ for all t , and (iii) there exists some time T_0 such that $u_t = 0$ for all $t \geq T_0$.

Proof: To prove (i) and (ii) jointly, we will proceed by induction on t . The base case holds by the assumption that we have access to an initial u_0 for which $y_1 \geq y_{\min}$. Assume

for the inductive hypothesis that $0 \leq u_{t-1} \leq \dots \leq u_0$ and $y_t \geq y_{\min}$. First, observe that since $y_t - y_{\min} \geq 0$, we have

$$u_t \leq \max\{0, u_{t-1}\} = u_{t-1}$$

showing the inductive step for (i). Second, by Lemma 1,

$$\begin{aligned} y_{t+1} &\geq y_{\min} + (y_{t+1}^{\text{nat}} - y_t^{\text{nat}}) - \sum_{k=1}^t g(k)(u_{t-k-1} - u_{t-k}) \\ &\geq y_{\min} + (y_{t+1}^{\text{nat}} - y_t^{\text{nat}}) + g(t)u_0 \\ &\geq y_{\min} + \delta/t \end{aligned}$$

where the second-to-last inequality follows since $g(k) \leq 0$ and $u_{t-k-1} - u_{t-k} \geq 0$ for all $1 \leq k \leq t-1$. The last inequality follows by our assumption $y_{t+1}^{\text{nat}} \geq y_t^{\text{nat}} - g(t)u_0 + \delta/t$. This shows the inductive step for (ii). Therefore, by induction, (i) and (ii) are true.

Finally, by the calculation above, we know that u_t is decreasing and, further, that it can be expressed as

$$u_t = \max\left\{0, u_0 - K_+ \delta \cdot \sum_{k=1}^t \frac{1}{k}\right\}$$

which, since $\sum_{k=1}^t 1/k \geq \ln t$, implies $u_t = 0$ for all $t \geq \exp\left\{\frac{u_0}{K_+ \delta}\right\}$, proving (iii). ■

VI. NUMERICAL SIMULATIONS

In this section, we present simulations demonstrating the behavior of the proposed protocols and standard (non-adaptive) tapering approaches. Due to a lack of available simulators/open source data, we rely on synthetic impulse response functions that we model based on broad properties of immediate response and allostatic adaptation.

a) Task definition: We consider the problem of tapering in the three opponent processes from Figure 3, starting from the allostatic state reached after taking dose $u = 1$ for $T_{\text{mit}} = 60$ timesteps. For each of these opponent processes, we are interested in the ability of a protocol to perform well on average for a population with different constraints. More concretely, for each setting, we consider ‘treating’ a population of N units over a horizon T_{taper} with y_{\min} distributed uniformly between the allostatic equilibrium and a lower baseline. Concretely, we take $y_{\min}^A \sim \text{Unif}(-2, 0)$, $y_{\min}^B \sim \text{Unif}(-1, 1)$, $y_{\min}^C \sim \text{Unif}(-4, -2)$ and $T_{\text{taper}}^A = T_{\text{taper}}^B = 60$, $T_{\text{taper}}^C = 14$. For all setups, we also inject random $\text{Unif}(-0.25, 0.25)$ noise to the observations to accommodate random disturbances from the underlying dynamics.

For a given individual and a given tapering protocol, we measure the average dose taken and the average constraint violation over the tapering horizon T_{taper} . We then average these two metrics over the N units to obtain the average performance of a protocol over a population.

b) Tapering protocol specification: As baselines, we consider the linear ($u_t = u_0 - \alpha t$) and the exponential ($u_t = \alpha^t u_0$) dose decay protocols and vary the decay rate α to get the trade-off curve between our two performance metrics for each opponent process when evaluated over units with varied constraints.

We compare against the integral controller of Section V for which we can directly input the patient constraint y_{\min} . The main obstacle is that we do not know the exact value of $g(0)$. However, since this is the immediate response, we can assume access to a course range. For our evaluation, we take the lower and upper bounds to be 50% and 150% of the true response. This directly corresponds to conservative settings for K_+ and K_- equal to $2g(0)^{-1}/3$ and $2g(0)^{-1}$. In the Appendix, we provide an ablation over the coarseness of our range specification, observing that our method can handle wide ranges. As discussed in Section V, our method can be instantiated with a padded $y_{\min}^\delta = y_{\min} + \delta$ to provide stronger constraint violation guarantees. Since this is the one hyperparameter that might affect practice, we also vary it in our experiments.

Finally, we compare to the optimal tapering protocol given by (3), for which we use the full model g and the true underlying natural progression to compute the maximally greedy dose. That is, we compare against the best possible control scheme had we known the true signal y_{t+1}^{nat} in advance. We also tested the cold turkey protocol ($u_t \equiv 0$), but its performance was too poor to be included.




Fig. 4. Average constraint violation (x-axis) against average dose (y-axis). Each point is averaged over $N = 100$ units.

c) Results.: In Figure 4, we see that the integral control law dominates the baselines considered, as the curves obtained by sweeping over the various parameters do not intersect. In the Appendix, we provide further experiments covering the robustness of the method to even more extreme misspecification of $g(0)$, as well as behavior in terms of a potential additional metric of interest: the percentage of units that successfully taper within the specified time frame.

Finally, there is a non-negligible performance gap between the integral and the optimal tapering protocols discussed in this paper, which indicates an avenue for future research.

There remain several control-theoretic considerations in tapering protocol design. For example, the desirable guarantee of monotonic dose decrease was induced by a discretization such that the full positive effect occurs within one timestep. How can we find the appropriate time window without excessive exploration? Can we develop methods for estimating the switching point of the process and discretizing such that the induced discrete system has $\tau_0 = 1$?

Additionally, numerous practical, clinical, and ethical considerations must be considered for the downstream deployment of the proposed methods. For example, neither the optimal policy nor our default integral control policy decrease the dose monotonically. For some individuals, this could result in recommending increasing dosages, which is undesirable in most clinical settings. A translation to practice may necessarily add the constraint that a dose never exceeds some level. And this may mean that desired levels of well-being are not achievable for some patients.

To our knowledge, this work is the first formalization of tapering as an optimal control problem. This formalism leads to a concrete approach that departs from existing literature on tapering and prescribes adaptivity. Excitingly, the proposed protocols may be able to address some of the limitations of current approaches while still providing simple, explainable rules. We hope that the straightforward protocols are amenable to undergoing human subject validation in future studies.

REFERENCES

- [1] R. L. Solomon, "The opponent-process theory of acquired motivation: the costs of pleasure and the benefits of pain." *American psychologist*, vol. 35, no. 8, p. 691, 1980.
- [2] I. Mujika and S. Padilla, "Scientific bases for precompetition tapering strategies," *Medicine & Science in Sports & Exercise*, vol. 35, no. 7, pp. 1182–1187, 2003.
- [3] M. A. Horowitz, R. M. Murray, and D. Taylor, "Tapering antipsychotic treatment," *JAMA psychiatry*, vol. 78, no. 2, pp. 125–126, 2021.
- [4] M. A. Horowitz, S. Jauhar, S. Natesan, R. M. Murray, and D. Taylor, "A method for tapering antipsychotic treatment that may minimize the risk of relapse," *Schizophrenia Bulletin*, vol. 47, no. 4, pp. 1116–1129, 2021.
- [5] C. Berna, R. J. Kulich, and J. P. Rathmell, "Tapering long-term opioid therapy in chronic noncancer pain: evidence and recommendations for everyday practice," in *Mayo clinic proceedings*, vol. 90, no. 6. Elsevier, 2015, pp. 828–842.
- [6] J. J. Fenton, A. L. Agnoli, G. Xing, L. Hang, A. E. Altan, D. J. Tancredi, A. Jerant, and E. Magnan, "Trends and rapidity of dose tapering among patients prescribed long-term opioid therapy, 2008–2017," *JAMA network open*, vol. 2, no. 11, pp. e1916271–e1916271, 2019.
- [7] A. Agnoli, G. Xing, D. J. Tancredi, E. Magnan, A. Jerant, and J. J. Fenton, "Association of dose tapering with overdose or mental health crisis among patients prescribed long-term opioids," *Jama*, vol. 326, no. 5, pp. 411–419, 2021.
- [8] S. G. Henry, D. A. Paterniti, B. Feng, A.-M. Iosif, R. L. Kravitz, G. Weinberg, P. Cowan, and S. Verba, "Patients' experience with opioid tapering: a conceptual model with recommendations for clinicians," *The Journal of Pain*, vol. 20, no. 2, pp. 181–191, 2019.
- [9] H. Kalant, "Biological models of alcohol tolerance and physical dependence," *Alcohol Intoxication and Withdrawal I: Experimental Studies*, pp. 3–14, 1973.
- [10] G. F. Koob and M. Le Moal, "Drug addiction, dysregulation of reward, and allostasis," *Neuropsychopharmacology*, vol. 24, no. 2, pp. 97–129, 2001.

- [11] G. F. Koob, "The dark side of emotion: the addiction perspective," *European journal of pharmacology*, vol. 753, pp. 73–87, 2015.
- [12] D. R. Wesson and W. Ling, "The clinical opiate withdrawal scale (cows)," *Journal of psychoactive drugs*, vol. 35, no. 2, pp. 253–259, 2003.
- [13] L. Handelsman, K. J. Cochrane, M. J. Aronson, R. Ness, K. J. Rubinstein, and P. D. Kanof, "Two new rating scales for opiate withdrawal," *The American journal of drug and alcohol abuse*, vol. 13, no. 3, pp. 293–308, 1987.
- [14] J. F. Rosenbaum, M. Fava, S. L. Hoog, R. C. Ascroft, and W. B. Krebs, "Selective serotonin reuptake inhibitor discontinuation syndrome: a randomized clinical trial," *Biological psychiatry*, vol. 44, no. 2, pp. 77–87, 1998.
- [15] K. Kroenke and R. L. Spitzer, "The phq-9: a new depression diagnostic and severity measure," pp. 509–515, 2002.
- [16] M. Mahdavi, R. Jin, and T. Yang, "Trading regret for efficiency: online convex optimization with long term constraints," pp. 2503–2528, 2012.
- [17] S. Mannor and J. N. Tsitsiklis, "Online learning with constraints," *Lecture notes in computer science*, vol. 4005, p. 529, 2006.

APPENDIX

a) *Proof of Lemma 1:* We first derive an alternate expression of the (possibly negative) dose u_t^* that leads to y_{\min} . Concretely, using the same derivation in Proposition 1 of the maximally greedy protocol (3), we have that

$$u_t^* = \frac{y_{\min} - y_{t+1}^{\text{nat}} - \sum_{k=1}^t g(k)u_{t-k}}{g(0)}$$

We aim to express this in terms of the current set-point error $y_t - y_{\min}$ and u_{t-1} , so we add and subtract y_t^{nat} and $\sum_{k=1}^{t-1} u_{t-k-1}$ (which generate $y_t - g(0)u_{t-1}$) to obtain:

$$\begin{aligned} u_t^* &= g(0)^{-1} (y_{\min} - y_t^{\text{nat}} + (y_t^{\text{nat}} - y_{t+1}^{\text{nat}}) \\ &\quad - \sum_{k=1}^{t-1} g(k)u_{t-k-1} - \sum_{k=1}^t g(k)(u_{t-k} - u_{t-k-1})) \end{aligned}$$

where we take $u_{-1} = 0$. Since $y_t^{\text{nat}} + \sum_{k=1}^{t-1} g(k)u_{t-k-1} = y_t - g(0)u_{t-1}$, we can rewrite the above as

$$\begin{aligned} u_t^* &= u_{t-1} - g(0)^{-1} (y_t - y_{\min}) - g(0)^{-1} (y_{t+1}^{\text{nat}} - y_t^{\text{nat}}) \\ &\quad + g(0)^{-1} \left(\sum_{k=1}^t g(k)(u_{t-k-1} - u_{t-k}) \right) \end{aligned}$$

Since $K_+ \leq g(0)^{-1}$ and $K_- \geq g(0)^{-1}$, we have that

$$-g(0)^{-1} (y_t - y_{\min}) \leq -K_+ (y_t - y_{\min})_+ - K_- (y_t - y_{\min})_-$$

And therefore, if u_t is given by (5), then

$$\begin{aligned} u_t^* &\leq u_t - g(0)^{-1} (y_{t+1}^{\text{nat}} - y_t^{\text{nat}}) \\ &\quad + g(0)^{-1} \left(\sum_{k=1}^t g(k)(u_{t-k-1} - u_{t-k}) \right) \end{aligned}$$

Multiplying both sides by $g(0)$ and then adding $\sum_{k=1}^t g(k)u_{t-k} + y_{t+1}^{\text{nat}}$, we have

$$\begin{aligned} &y_{t+1}^{\text{nat}} + g(0)u_t^* + \sum_{k=1}^t g(k)u_{t-k} \\ &\leq y_{t+1}^{\text{nat}} + \sum_{k=0}^t g(k)u_{t-k} - (y_{t+1}^{\text{nat}} - y_t^{\text{nat}}) \\ &\quad + \sum_{k=1}^t g(k)(u_{t-k-1} - u_{t-k}) \end{aligned}$$

Since by construction the LHS is equal to exactly y_{\min} and by the dynamics equation $y_{t+1}^{\text{nat}} + \sum_{k=0}^t g(k)u_{t-k} = y_t$, rearranging the above yields:

$$y_{\min} + (y_{t+1}^{\text{nat}} - y_t^{\text{nat}}) - \sum_{k=1}^t g(k)(u_{t-k-1} - u_{t-k}) \leq y_t$$

which is precisely the stated inequality.

b) *Generalization to non-linear opponent processes:*

Our definition of (1) implies a linear contribution of the dose to the observed well-being metric. One may wonder if we can accommodate the case where the opponent process varies with u , i.e. y_t evolves according to the following generalized version of (1):

$$y_{t+1} = \sum_{k=0}^t g(k, u_{t-k}) + y_{t+1}^{\text{nat}} \quad (6)$$

In this subsection, we show we can, under a corresponding generalization of well-behavedness from Definition 2. The rate of change of each $g(t, \cdot)$ with respect to the dose u will play an important role, so we introduce the following two shorthand notations:

$$\partial_u^-(t) \doteq \inf_u \partial_u g(t, u) \quad \partial_u^+(t) \doteq \sup_u \partial_u g(t, u)$$

Definition 4 (Generalized Opponent Process): We say that g is a generalized opponent process if there exists a time τ_0 such that:

- 1) $\inf_u g(\tau, u) \geq 0$ and $\partial_u^-(t) \geq 0$ when $\tau < \tau_0$,
- 2) $\sup_u g(\tau, u) \leq 0$ and $\partial_u^+(t) \leq 0$ when $\tau \geq \tau_0$.

Definition 5 (Generalized LPOP): A generalized opponent process g is a generalized linearly progressing opponent process (G-LPOP) if there exists an $\alpha \in (0, 1)$ such that:

- 1) $\partial_u^+(t+1) \leq \alpha \cdot \partial_u^-(t)$ when $t < \tau_0 - 1$,
- 2) $\partial_u^-(t+1) \geq \alpha \cdot \partial_u^+(t)$ when $t \geq \tau_0$.

In this case the proof of Theorem 1 changes a bit, relying on Mean Value Theorem. We give the statement and proof in Theorem 3 below.

Theorem 3: For any G-LPOP g , taking the maximally greedy dose at every time maintains $y_t \geq y_{\min}$ for all t with minimal cumulative dose.

Proof: Consider a dosing sequence u_t^* that maintains $y_t^* \geq y_{\min}$ for all t with minimal cumulative dose. Assume it is not maximally greedy (otherwise we would be done). This means that there exists t_0 and $\epsilon > 0$ such that $u_{t_0}^* - \epsilon$ would produce a new $y_{t_0+1}^* \geq y_{\min}$. Consider the modified dosing schedule

$$u'_t = \begin{cases} u_{t_0}^* - \epsilon, & t = t_0 \\ u_{t_0+1}^* + \alpha\epsilon, & t = t_0 + 1 \\ u_t^* & \text{otherwise} \end{cases}$$

and let y'_t be the resulting observations. Since $u'_t = u_t^*$ for $t < t_0$, we have that $y'_t = y_t^* \geq y_{\min}$ up to t_0 . Plugging in the effect of the dose modifications at t_0 and $t_0 + 1$ into

the generalized dynamics equation (6), by the Mean Value Theorem, we get:

$$\begin{aligned} y'_{t+1} &= y_{t+1}^* - g(t - t_0, u_{t_0}^*) + g(t - t_0, u_{t_0}^* - \epsilon) \\ &\quad - g(t - t_0 - 1, u_{t_0+1}^*) + g(t - t_0 - 1, u_{t_0+1}^* + \alpha\epsilon) \\ &= y_{t+1}^* - \epsilon \cdot \partial_u g(t - t_0, u_{t_0}^* - \delta_0\epsilon) \\ &\quad + \alpha\epsilon \cdot \partial_u g(t - t_0 - 1, u_{t_0+1}^* + \delta_1\epsilon) \end{aligned}$$

for some $\delta_0 \in (0, 1), \delta_1 \in (0, \alpha)$. Since g is a G-LPOP, this final expression is greater than or equal to y_{t+1}^* and hence is also greater than or equal to y_{\min} . Since this holds for all $t > t_0$, we have that y'_t satisfies the constraints for all t . To finalize, note that y'_t is achieved with smaller cumulative dose since

$$\sum_{t=1}^T u'_t = \sum_{t=1}^T u_t^* - \epsilon + \alpha\epsilon < \sum_{t=1}^T u_t^*,$$

This is a contradiction, completing the proof. ■

Instantiating this in the form of protocol (2) can be done by solving for the dose which exactly matches the setpoint:

Proposition 3: For any G-LPOP g , assume we know a lower bound $y_t^{\text{nat,lb}}$ such that $y_t^{\text{nat,lb}} \leq y_t^{\text{nat}}$ for all t . Then, for any instance of y_t^{nat} that is greater than the prescribed lower bound, the dosing sequence which takes u_t be the solution to

$$g(0, u) = \max \left\{ 0, y_{\min} - y_{t+1}^{\text{nat,lb}} - \sum_{k=1}^t g(k, u_{t-k}) \right\} \quad (7)$$

maintains $y_t \geq y_{\min}$ for all t with minimal cumulative dose.

Remark 1: Since $g(0, u)$ is increasing in u , an ϵ -accurate solution to (7) can be computed in $\mathcal{O}(\log \frac{1}{\epsilon})$ steps.

The above effectively shows that optimal tapering is equivalent to setpoint matching under non-linear g as well. While it is not as clean to derive precise guarantees for the integral controller as in the main text, the result of Theorem 3, combined with the robustness to $g(0)$ errors observed experimentally, strongly motivates the same approach in this more general setting.

c) Additional Experiment (Ablation over $g(0)$ range coarseness): Figure 5 demonstrates that the integral controller can support a wide misspecification on the range of $g(0)$. In fact, the experiments suggest an overly conservative upper bound on $g(0)$ may improve performance under an optimal choice of δ .

d) Additional Experiment (Number of subjects fully tapered): In Figure 6 we plot the percentage of units the proposed protocol fully tapered within T_{taper} steps for varying values of δ . We see that our method fully tapers a portion of the units within some pre-specified timeline with minimal constraint violation. This is in contrast with fixed methods which induce a 0% vs. 100% tapering success step function given a timeframe T_{taper} . In all settings used, non-adaptive methods would have to incur significant constraint violations to fully taper within the pre-specified timeframe.




Fig. 5. Average constraint violation (x-axis) against average dose (y-axis) for different ranges of lower and upper bounds on $g(0)$. A label of (p_1, p_2) in the legend corresponds to setting K to $(p_1 g(0))^{-1}$ and K_+ to $(p_2 g(0))^{-1}$.




Fig. 6. Average constraint violation (x-axis) against the percentage of units fully tapered (y-axis).