

Control Systems Engineering for Understanding and Optimizing Smoking Cessation Interventions*

Kevin P. Timms¹, Daniel E. Rivera², Linda M. Collins³, and Megan E. Piper⁴

Abstract—Cigarette smoking remains a major public health issue. Despite a variety of treatment options, existing intervention protocols intended to support attempts to quit smoking have low success rates. An emerging treatment framework, referred to as adaptive interventions in behavioral health, addresses the chronic, relapsing nature of behavioral health disorders by tailoring the composition and dosage of intervention components to an individual’s changing needs over time. An important component of a rapid and effective adaptive smoking intervention is an understanding of the behavior change relationships that govern smoking behavior and an understanding of intervention components’ dynamic effects on these behavioral relationships. As traditional behavior models are static in nature, they cannot act as an effective basis for adaptive intervention design. In this article, behavioral data collected daily in a smoking cessation clinical trial is used in development of a dynamical systems model that describes smoking behavior change during cessation as a self-regulatory process. Drawing from control engineering principles, empirical models of smoking behavior are constructed to reflect this behavioral mechanism and help elucidate the case for a control-oriented approach to smoking intervention design.

I. INTRODUCTION

Behavioral interventions describe a class of preventive or therapeutic measures intended to reduce unhealthy behaviors or promote healthy ones [1]. Addressing a wide variety of public health issues such as substance abuse, obesity, and disease screening, these programs are pharmacological, behavioral (e.g., clinician counseling), or community-based (e.g., public smoking regulations) in nature [1], [2]. Due to a multitude of complicating factors such as multi-level and interacting sources of inter- and intra-individual variability (e.g., gene-environment interactions), effective design of behavioral interventions poses a major challenge [3], [4]; such challenges are compounded by the array of medical, ethical, and resource limitations demanded of interventions, as well as their experimental design and evaluation.

*This work was supported by an award from the American Heart Association, the Office of Behavioral and Social Sciences Research (OBSSR) of the National Institutes of Health (NIH) and the National Institute on Drug Abuse (NIDA) through grants K25 DA021173, R21 DA024266, and P50 DA10075.

¹Biological Design Program, Arizona State University, Tempe, AZ ktimms@asu.edu

²Control Systems Engineering Laboratory, School for Engineering of Matter, Transport and Energy, Arizona State University, Tempe, AZ daniel.rivera@asu.edu

³The Methodology Center and Department of Human Development and Family Studies, Pennsylvania State University, State College, PA lmcollins@psu.edu

⁴Department of Medicine, Center for Tobacco Research and Intervention, University of Wisconsin, Madison, WI MEP@ctri.medicine.wisc.edu

Behavioral interventions designed to promote cigarette smoking cessation are a major public health concern: in the U.S. alone, approximately 440,000 premature deaths and \$157 billion in economic loss are attributed to smoking annually [5]. A significant obstacle in addressing this public health issue, though, is the limited utility of existing intervention protocols. Counseling alone, for example, has a reported successful quit rate below 15% and pharmaco-therapies (e.g., nicotine gum) have individual success rates below 35% [2]. These existing interventions are “fixed,” meaning all patients receive the same treatment protocol, which does not vary over the course of the intervention [3].

As an alternative, *adaptive interventions* have emerged as a promising strategy to address the chronic, relapsing nature of cigarette smoking [3], [2]. In adaptive interventions, treatments are tailored to individual patients over time and are defined according to *decision rules*, which specify intervention composition and dosage based on *tailoring variables*, such as a real-time patient-reported behaviors [3].

Traditionally, behavioral science models are static in nature. These models’ inability to describe behavioral and treatment dynamics means they cannot effectively inform adaptive intervention designs [6]. Toward a better understanding of intervention dynamics, computerized technologies offer a cost-effective means to collect *intensive longitudinal data* (ILD)—behavioral data collected at short intervals over time [6]. Figure 1 features such ILD, collected from participants in a University of Wisconsin smoking cessation clinical trial through personal digital assistants (PDAs) nightly. Specifically, *Craving* and *Cigsmked* (the number of cigarettes smoked per day), are shown for four data sets, averages for two of the study’s treatment groups and for two single participants from those groups [7]. As seen in Figure 1, the ability of ILD to capture behavioral dynamics offers the opportunity to examine smoking behavior change from a dynamical systems perspective [4], [8], [6]. From an engineering perspective, ILD used with system identification modeling techniques makes a connection between behavioral healthcare and control-guided intervention design [8], [9]. For example, empirical models that describe dynamical features, such as *Craving*’s inverse response or *Cigsmked*’s dramatic reduction immediately following quit and possible post-quit resumption in Fig. 1, could act as the plant model ($\tilde{P}(s)$) in a smoking intervention algorithm [8]. More specifically, behavioral models developed from ILD measured daily lay the groundwork for a “just-in-time” adaptive intervention. In behavioral science settings, this describes adaptation implemented during the intervention—but with minutes to days

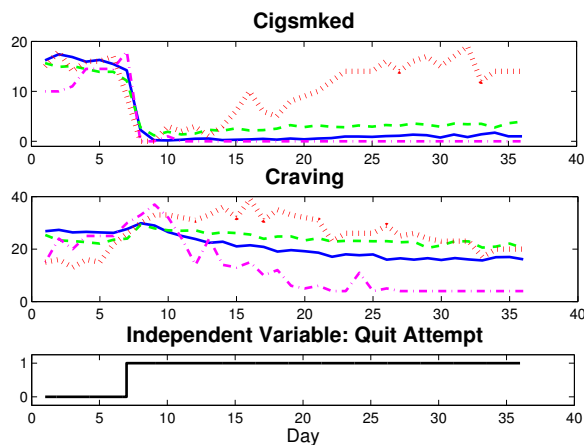


Fig. 1. *Craving* and *Cigsmked* data collected daily in a UW-CTRI smoking cessation clinical trial. Four sets of signals are shown: averaged data for a group receiving bupropion and counseling as treatments (AC group; solid blue), averaged data for a group receiving neither drug nor counseling (PNc group; dashed green), a single subject from the AC group (dash-dot magenta), and a single subject from the PNc group (dotted red).

of lag between real-time changes in tailoring variables and assignment of intervention adjustments [6].

One of the most significant ways in which behavioral science can benefit from a control-oriented interpretation is in analysis of feedback and feedforward processes [8]. Self-regulation is a theory postulated in behavioral science [10], and has been theorized as a mechanism governing cigarette smoking maintenance. Specifically, the Nicotine Regulation Model and similar hypotheses suggest that smoking is done to maintain a blood nicotine or urge set point [11], [12]. Given that one of the greatest advantages of ILD is an opportunity to identify and describe self-regulatory and self-exciting relationships among behavioral variables [4], control systems engineering offers a mature methodology to comprehensively and parsimoniously characterize smoking cessation behavior change [8], [9].

Altogether, control systems engineering offers a comprehensive platform to address the feedback phenomena that enters at two levels in the understanding smoking behavior change and subsequent intervention design. Specifically, models of self-regulatory smoking cessation behavior change could be used in conjunction with controller design principles to produce an intervention algorithm that systematically defines a patient-specific treatment plan based on a patient's previous behavior, changing needs, anticipated disturbances, and more. A significant advantage of such a control-theory-based-design is the opportunity to optimize for intervention efficacy and speed, while adhering to resource, medical, and other operational constraints [1], [8].

This paper is organized as follows: a University of Wisconsin Center for Tobacco Research and Intervention (UW-CTRI) smoking cessation clinical trial is first outlined [7]. Next, a model of smoking behavior change that features a feedback loop is presented and the UW-CTRI data is used to produce empirical smoking behavior change models. The self-regulation model is then discussed in the context of an

intervention. Finally, conclusions are drawn and recommendations for future work are made.

II. SMOKING CESSATION CLINICAL TRIAL

ILD from a Transdisciplinary Tobacco Use Research Center-funded UW-CTRI study provide an experimental basis for evaluation of critical relationships among smoking behaviors. This clinical trial examined counseling and bupropion SR as smoking cessation aids: 101 participants received active drug and counseling (“AC” group), 100 received active drug and no counseling (“ANc”), 101 received placebo drug and counseling (“PC”), and 101 received placebo drug and no counseling (“PNc”). Those receiving active bupropion, which has been shown to alleviate withdrawal symptoms, including craving, took 150 mg starting one week prior to the planned quit date and 300 mg starting four days prior to the planned quit, allowing the drug to build up in the participants’ systems. The counseling protocol entailed eight counseling sessions that covered topics such as coping, motivation, and preventing relapse [7].

ILD was collected nightly through personal digital assistants (PDAs) for all participants in the form of Evening Reports (i.e., questionnaires); these self-reports recorded average levels of a variable over the day and actual smoking behaviors (e.g., *Cigsmked*, the total number of cigarettes smoked that day).

This article focuses on the relationship between *Craving*—the sum of *Urge*, *Cigonmind*, *Thinksmk*, and *Bother* questions, each measured on a 10-point Likert scale (1-11, No!!...Yes!!)—and *Cigsmked*. Two sets of group average *Craving* and *Cigsmked* signals are depicted in Fig. 2 for one week prior to quit through the first four post-quit weeks (solid blue, AC; dash-dot magenta, PNc). Fig. 2 indicates similar dynamic features for the group averages: *Craving* features distinct inverse response upon initiation of a quit attempt, and *Cigsmked* features a large, fast reduction in smoking, followed by a relatively small and slow resumption of smoking. Fig. 3 features data for one subject from the AC group (solid blue, with eight days of data imputed) and one from the PNc group (dash-dot magenta, seven days of missing data imputed). Although both single subject data sets are noisier than the group-average data shown, the AC subject’s *Craving* signal clearly features inverse response and *Cigsmked* has no resumption, suggesting a successful quit attempt. Conversely, the PNc subject features a dramatic reduction in *Cigsmked* and an increase in *Craving* upon quitting, but appears to fully resume smoking, as *Craving* and *Cigsmked* settle to approximately pre-quit levels.

III. SMOKING BEHAVIOR SELF-REGULATION MODEL

Behavioral science’s particular interest in the concept of self-regulation for describing smoking behavior warrants a control engineering interpretation of smoking cessation [11], [12], [9]. Furthermore, estimation of such models is motivated by the fact that dynamic models describing smoking cessation as a process featuring mediation—a multivariable

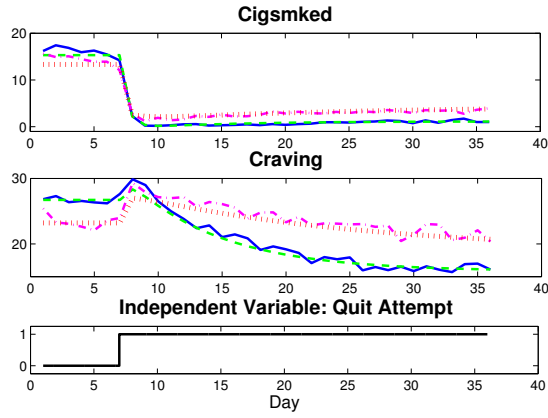


Fig. 2. Plot of *Craving* and *Cigsmked* group average data (solid blue, AC; dash-dot magenta, PNc) and corresponding dynamical systems models (dashed green, AC; dotted red, PNc).

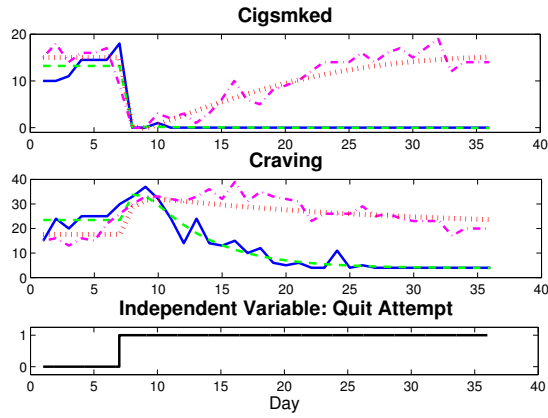


Fig. 3. Plot of *Craving* and *Cigsmked* data and dynamical systems models for two participants from the UW-CTRI study (solid blue, AC subject data; dashed green, AC subject model; dash-dot magenta, PNc subject data; dotted red, PNc subject model).

causal relationship central to mechanistic modeling in the social sciences—do not fully describe a *Craving-Cigsmked* interrelationship captured in the UW-CTRI ILD [9].

The solid lines in Fig. 4 feature a block diagram depicting a novel mechanism that describes smoking cessation behavior change as a feedback process. In this model, *Craving* is directly related to *Cigsmked* by $P(s)$, and *Cigsmked* itself is composed of two pathways, a quit pathway and a feedback pathway. The quit pathway’s contribution is the output of $P_d(s)$, which accepts a unit step input, $Quit(s)$, that corresponds to the transition from not attempting to quit smoking to attempting to quit. The feedback pathway’s contribution is the output of $C(s)$, a craving self-regulator, which accepts e —the deviation between a craving set point, r , and the time-varying *Craving* signal—as the input. This proposed model describes an intuitive relationship: smoking is governed by the difference between an individual’s biologically and/or psychologically-determined perceived need to smoke (r), and time-varying craving for cigarettes, which itself is a result of smoking; an attempt to quit smoking is effectively a disturbance on this regulatory process.

The closed-loop transfer functions corresponding to the solid lines in Fig. 4 are,




Fig. 4. Block diagram depicting a smoking behavior change mechanism in which a *Craving-Cigsmked* interrelationship is described by self-regulation (solid); a hypothetical treatment mechanism (dashed).

$$Craving = \left(\frac{PC}{1+PC} \right) r + \left(\frac{PP_d}{1+PC} \right) Quit \quad (1)$$

$$Cigsmked = \left(\frac{C}{1+PC} \right) r + \left(\frac{P_d}{1+PC} \right) Quit \quad (2)$$

With the UW-CTRI ILD and (1) and (2), empirical model development becomes a closed-loop identification problem. For this, $P(s)$ is estimated as a single-input-single-output system with *Cigsmked* as the input and *Craving* as the output. $P_d(s)$ and $C(s)$ are estimated as a multi-input-single-output system with $Quit$ and $e (=r - Craving)$ as inputs and *Cigsmked* as the output.

In evaluating the proposed mechanism, a nomothetic (i.e., group-focused) perspective is employed to develop the empirical models using group average ILD from the UW-CTRI study. Doing so can provide generalizable insight into behavioral relationships. $P(s)$, $P_d(s)$, and $C(s)$ are estimated using the PEM functionality and *process models* notation in Matlab’s System Identification Toolbox. In this analysis, the goodness-of-fit metric (0-100% scale) is calculated as,

$$Fit [\%] = 100 \left(1 - \frac{\|y(t) - \hat{y}(t)\|_2}{\|y(t) - \bar{y}\|_2} \right) \quad (3)$$

where $y(t)$ is the data, \bar{y} is the mean value of the 36-day data signal to which the model is fit, and $\hat{y}(t)$ is the simulated output. In the iterative estimation process used here, transfer function structures were varied and the resulting candidate models simulated. Validity of the resulting empirical models were evaluated for predictive ability, parameter realizability, and transfer function parsimony. The appropriate r was determined during estimation, as several constant r values and linear functions were tested.

The results of this estimation are found in Table I and the associated simulations for the AC and PNc groups are found in Fig. 2 (dashed green, AC; dotted red, PNc). Although five distinct stages of smoking cessation have been theorized [13], which may suggest r may vary over time, the pre-quit baseline *Craving* level appears to be the most appropriate r over the post-quit time period modeled. The appropriate transfer function structures that correspond to the parameters estimated using the group average ILD in Table I were determined to be,

$$P(s) = \frac{K_1(\tau_a s + 1)}{(\tau_1 s + 1)} \quad (4)$$

$$P_d(s) = K_d \quad (5)$$

$$C(s) = \frac{K_c}{(\tau_{c1} s + 1)} \quad (6)$$

TABLE I

FEEDBACK MODEL PARAMETER ESTIMATES AND GOODNESS-OF-FITS.

Treatment: Data Set:	AC Avg	ANc Avg	PC Avg	PNC Avg	AC Sgl	PNC Sgl
Craving Fit [%]	87.32	77.65	77.51	62.25	66.90	57.59
Cigsmked Fit [%]	89.16	83.03	91.44	84.12	77.09	62.99
Craving Set. T.	35.69	35.91	35.82	35.90	26.34	33.87
Cigsmked Set. T.	34.56	35.26	35.60	35.29	10.64	33.86
$P(s) K_1$	0.77	0.74	0.50	0.52	1.57	-2.21
$P(s) \tau_{a1}$	-1.99	-3.76	-14.34	-21.90	-3.05	3.45
$P(s) \tau_1$	8.22	14.23	18.70	26.75	6.88	10.76
$C(s) K_c$	0.08	0.23	0.20	0.30	-0.01	-6.25
$C(s) \tau_c$	4.59	2.89	0.42	1.89	1.29	95.53
$P_d(s) K_d$	-15.01	-13.13	-13.50	-10.24	-13.05	-15.99

As documented in Table I, high goodness-of-fit values are achieved for the four groups, which all feature comparable dynamic phenomena. This suggests that the self-regulatory process depicted as the solid lines in Fig. 4 appropriately represents a fundamental relationship between craving and daily smoking. This behavior change mechanism is supported by the fact that models where *Craving* and *Cigsmked* are reversed have no predictive ability. Interestingly, we find that $P_d(s)$ models the dramatic reduction in daily smoking upon quitting, while the $C(s)$ pathway models the resumption in smoking. More importantly, the $C(s)$ transfer function structure indicates that across a population, a smoker's craving self-regulator behaves as a Proportional-with-Filter controller. Furthermore, it is known in dynamical systems that a model with a zero term can come from two parallel first order models. For the group average models, the first subprocess, $P_1(s)$, would have a negative gain and a smaller speed of response compared to that for a second subprocess, $P_2(s)$, which would have a positive gain. *Craving's* initial increase would reflect the fast contribution of the negative-gained function as it accepts the large decrease in *Cigsmked* upon quitting, while *Craving's* settling to below pre-quit levels would reflect the slow contribution of the positive-gained function as it accepts the increasing *Cigsmked* values (resumption). The dual nature of this possible underlying $P(s)$ phenomena is appealing considering the two paired components of nicotine addiction (biochemical and conditioning processes), and the impulsive and executive neurological systems that manage immediate and delayed gratification motivations, respectively, which compete during smoking cessation [14].

Ideally, assignment and adjustment of intervention components would be based on patient-specific models of smoking

behavior change. This motivated development of empirical models for the two UW-CTRI single subject examples. Estimation and validation of these models involved the same procedure as before, but r was assumed to be the baseline craving level. The model structures that describe the group average dynamics adequately describe the single subject dynamics as well. The corresponding parameter estimates and simulations are in Table I and Fig. 3, respectively.

As the PNC subject's K_1 estimate is negative, the parallel subprocesses that may make up $P(s)$ would have a relationship different than that described for the group average models. Specifically, both subprocess gains would be negative, corresponding to the increase in *Craving* that results from the step-like decrease in *Cigsmked* on the quit date and the settling of *Craving* to approximately pre-quit values that results from the post-quit increase in *Cigsmked*. While still adequately described as a Proportional-with-Filter controller, the PNC subject's $C(s)$ reflects this subject's significant resumption in daily smoking. The resumption in *Cigsmked* resulting from this subject's craving self-regulator highlights the chronic and relapsing nature of nicotine addiction, which may be effectively managed by interventions that are adaptive in nature [3], [8].

Examining Table I, the single subject models have lower goodness-of-fit values than their group average counterparts. These lower values represent a significant and expected challenge in single subject modeling; the reliability of single subject models is a function of the quality of source data, which may be noisy, require imputation of missing values, and more. The AC single subject's *Cigsmked* baseline exemplifies such concerns; the baseline appears to be non-stationary, but this is largely a consequence of the imputation of several consecutive days of missing data.

Although an adaptive smoking intervention algorithm would ideally be used in conjunction with accurate patient-specific behavior change models, identification of these models would require data collected during previous quit attempts. While design of a "patient-friendly" experiment could mitigate some single subject data quality issues, suggesting that a patient attempt to quit smoking for the purposes of reliable model estimation is largely impractical. However, a model of the smoking cessation process that reflects a behavioral predisposition to quit attempt failure can inform intervention design. The following expressions represent such a failed quit process for a hypothetical patient:

$$P(s) = \frac{-2.46(3.89s + 1)}{(15.05s + 1)} \quad (7)$$

$$P_d(s) = \frac{-20.29}{(1.77s + 1)} \quad (8)$$

$$C(s) = \frac{-20.07}{(90.92s + 1)} \quad (9)$$

In the context of the closed-loop smoking cessation process depicted as the solid lines in Figure 4, Equations (7), (8), and (9) represent a hypothetical patient's dynamics: the mathematical structures adhere to those that adequately

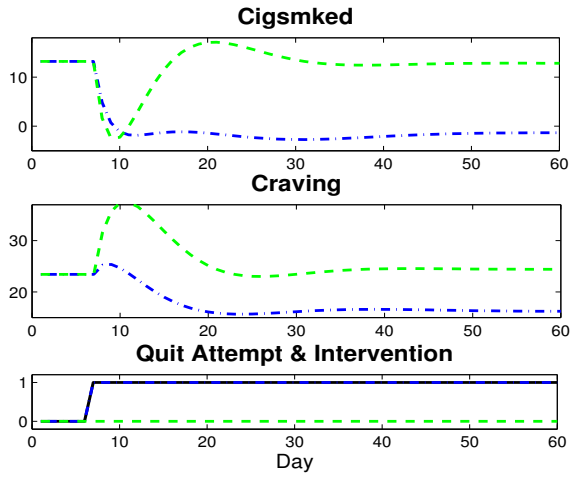


Fig. 5. Two months of simulated *Craving* and *Cigsmked* outputs for two cases with the quit attempt beginning on day 7 (coded as step, solid black): no intervention, dashed green; intervention, which is said to become effective on the quit date (coded as a step), dash-dot blue.

describe the clinical trial data, but the parameters reflect a behavioral relationship that inherently leads to failure of an attempt to quit smoking. Specifically, the parameters describe dynamics in which *Craving* increases upon a quit-induced reduction in *Cigsmked*, but both *Craving* and *Cigsmked* settle to approximately pre-quit levels. Selection of these parameters was guided by the apparent quit attempt failure in the previously described PNc single subject, although the PNc subject's steady-state *Craving* and *Cigsmked* models settle at values unequal to pre-quit levels.

IV. TOWARD A CLOSED-LOOP INTERVENTION

With an improved understanding of the fundamental smoking cessation behavioral mechanism, an analysis of medication effects is of interest. Here, a treatment scenario is modeled in which representative behavioral dynamics, i.e., model of a failed cessation attempt, are intervened upon to result in cessation success. This is a precursor to development of an optimal treatment strategy. In this section, one candidate pharmacological intervention mechanism that features dual models of action is examined.

A hypothetical pharmacological treatment is depicted as the dashed lines in Fig. 4. For this, $P_C(s)$ directly affects time-varying craving levels, which many existing smoking cessation medications are said to do, and $P_r(s)$ affects r , which represents the treatment interfering with addiction mechanisms [15]. Fig. 4's closed-loop transfer functions are,

$$Craving = \left(\frac{P_C + P_r PC}{1 + PC} \right) u_I + \left(\frac{PP_d}{1 + PC} \right) Quit \quad (10)$$

$$Cigsmked = \left(\frac{(P_r - P_C)C}{1 + PC} \right) u_I + \left(\frac{P_d}{1 + PC} \right) Quit \quad (11)$$

The treatment mechanism depicted in Fig. 4 is hypothetical, representing one possible relationship by which self-regulated smoking behavior change is affected by a pharmacological intervention. To examine this hypothetical intervention, simulations were conducted in Matlab's Simulink

environment where Equations (7), (8), and (9) represented the fundamental behavioral relationship. As the dynamic treatment mechanism discussed is not intended to reflect pharmacokinetics or pharmacodynamics of any specific pharmacological smoking intervention, selection of the $P_C(s)$ and $P_r(s)$ transfer function structures and parameters was guided by general knowledge of how a behavioral process predisposed to cessation failure could be altered for success, but were largely determined through trial-and-error. In identifying the intervention transfer function structures and parameters that lead to a more successful quit, u_I was set as a unit step occurring on the quit date; this corresponds to a fixed intervention protocol that becomes effective on the quit date. A candidate set of transfer functions promoting a successful quit are,

$$P_r(s) = \frac{-7.5}{(5s + 1)} \quad (12)$$

$$P_C(s) = \frac{-43}{(10s + 1)} \quad (13)$$

These transfer functions suggest the dual medication effects could independently (regardless of a quit attempt) facilitate a gradual decrease in r , an intuitive medication effect, and a decrease in *Craving*. Using the hypothetical patient's $P(s)$, $P_d(s)$, and $C(s)$ expressions and the $P_r(s)$ and $P_C(s)$ intervention models, Fig. 5 features 60 days of simulated *Craving* and *Cigsmked* results of this fixed intervention (intervention, dash-dot blue; intervention-free, dashed green).

Favorable behavioral dynamics are achieved in both *Craving* and *Cigsmked* in the simulation. Specifically, *Craving* decreases by approximately 30%, and the max peak is now 25.4, more than 10 points below the intervention-free simulation, and only about one point above the pre-quit level. Examining the *Cigsmked* simulation, this fixed intervention also appears to reach and maintain a successful quit.

Generally, these simulations illustrate that a fixed intervention protocol relying on a pharmacological smoking treatment that features dual modes of action can dynamically promote a successful quit within a behavioral system that otherwise leads to a failed cessation attempt.

The intervention mechanism presented may be suboptimal. First, the dynamic effects of existing pharmacological interventions have not been characterized according to the mechanism depicted in Fig. 4. Furthermore, inter-individual variability in the self-regulatory relationship (person-to-person variation in $P(s)$, $P_d(s)$, and $C(s)$ parameters), response to medication (person-to-person variation in $P_r(s)$ and $P_C(s)$ parameters), and unknown disturbances (exogenous variables, such as environmental smoking cues) suggest a fixed intervention similar to that simulated in Fig. 5 would be suboptimal, at least in terms of resource use. As identification of reliable patient-specific models that account for such variability is prohibitively challenging, the case for an adaptive intervention remains clear, as is the utility of controller design principles in optimizing such interventions. Specifically, controller design principles could be used in




Fig. 6. Proposed framework for an optimized adaptive smoking intervention.

conjunction with the self-regulatory behavioral model in a manner similar to that in Fig. 6. Design of such a controller could rely on medication to promote favorable signal shapes for both *Cigsmked* and *Craving*; while a *Cigsmked* = 0 set point would be prioritized, including a second set point for *Craving*, or treating this as an associated variable, could promote post-quit *Craving* reduction. Drawing from additional knowledge of relationships between behavioral, biological, and environmental variables (i.e., disturbances), a controller featuring both feedback and feedforward control is particularly appealing. For example, accounting for *Quit* as a forecasted measured disturbance may help mitigate negative behavioral effects of quitting; similarly, anticipating measured disturbances, such as an exogenous *Stress* signal, may also help manage the conditions promoting relapse. Altogether, a Hybrid Model Predictive Control (HMPC) strategy should be employed. Considering the human health focus of this problem, the ability of HMPC to manage the discrete nature of manipulated variables (i.e., pre-determined medication dosages) and hard- and soft-constraints makes it an advantageous intervention framework [16]. Furthermore, an HMPC design method could be used to provide insight into aspects of intervention optimality; notably, questions regarding necessary amounts of intervention components (e.g., medication dosage), duration of the overall treatment, and scheduling of alternate treatment components (e.g., combined medication and counseling) could be explored. Development of such an intervention algorithm could also inform design of novel clinical trials that more effectively measure and evaluate the dynamics and optimality of smoking cessation treatment strategies.

V. CONCLUSIONS AND RECOMMENDATIONS

In this paper, a control-oriented approach to the study of smoking cessation is employed. In doing so, the utility of such a perspective in this behavioral context is highlighted by use of ILD to produce parsimonious descriptions of self-regulation within smoking cessation. Models of this nature offer more comprehensive descriptions of smoking behavior as compared to traditional, static behavioral models. To develop an optimized smoking cessation intervention, a second level of feedback could be introduced to this problem. Specifically, an adaptive smoking intervention algorithm could be constructed to optimize and personalize cessation therapies. An algorithm of this nature could feature both feedback and feedforward components and would likely benefit from an HMPC framework.

REFERENCES

- [1] D.E. Rivera, M.D. Pew, and L.M. Collins. Using engineering control principles to inform the design of adaptive interventions: A conceptual introduction. *Drug and Alcohol Dependence*, 88(Supplement 2):S31–S40, 2007.
- [2] Tobacco Use and Dependence Guideline Panel. A clinical practice guideline for treating tobacco use and dependence: 2008 update. *American Journal of Preventive Medicine*, 35(2):158–176, 2008.
- [3] L.M. Collins, S.A. Murphy, and K.L. Bierman. A conceptual framework for adaptive preventive interventions. *Prevention Science*, 5(3):185–196, 2004.
- [4] L.M. Collins. Analysis of longitudinal data: The integration of theoretical model, temporal design, and statistical model. *Annual Review of Psychology*, 57(1):505–528, 2006.
- [5] P.R. Killeen. Markov model of smoking cessation. *PNAS*, 108 Supplement 3(37):15549–15556, 2011.
- [6] W.T. Riley, D.E. Rivera, A.A. Atienza, W. Nilsen, S.M. Allison, and R. Mermelstein. Health behavior models in the age of mobile interventions: Are our theories up to the task? *Translational Behavioral Medicine*, 1(1):53–71, 2011.
- [7] D.E. McCarthy, T.M. Piasecki, D.L. Lawrence, D.E. Jorenby, S. Shiffman, M.C. Fiore, and T.B. Baker. A randomized controlled clinical trial of bupropion SR and individual smoking cessation counseling. *Nicotine and Tobacco Research*, 10(4):717–729, 2008a.
- [8] D.E. Rivera. Optimized behavioral interventions: What does system identification and control systems engineering have to offer? *16th IFAC Symposium on System Identification, Brussels, Belgium*, pages 882–893, 2012.
- [9] K.P. Timms, D.E. Rivera, L.M. Collins, and M.E. Piper. System identification modeling of a smoking cessation intervention. *16th IFAC Symposium on System Identification, Brussels, Belgium*, pages 786–791, 2012.
- [10] C.S. Carver and M.F. Scheier. *On the Self-Regulation of Behavior*. Cambridge University Press, New York, New York, 1998.
- [11] W.F. Velicer, C.A. Redding, R.L. Richmond, J. Greeley, and W. Swift. A time series investigation of three nicotine regulation models. *Addictive Behaviors*, 17(4):325–345, 1992.
- [12] T.A. Walls and D.E. Rivera. Control engineering-based approaches to modeling substance abuse data. Washington, D.C., May 2009. 17th Annual Meeting of the Society for Prevention Research.
- [13] W.F. Velicer, S.L. Hughes, J.L. Fava, and J.O. Prochaska. An empirical typology of subjects within stage of change. *Addictive Behaviors*, 20(3):299–320, 1995.
- [14] W.K. Bickel, M.L. Miller, R. Yi, B.P. Kowal, D.M. Lindquist, and J.A. Pitcock. Behavioral and neuroeconomics of drug addiction: Competing neural systems and temporal discounting processes. *Drug and Alcohol Dependence*, 90(1):S85–S91, 2007.
- [15] R. West, C.L. Baker, J.C. Cappelleri, and A.G. Bushmakin. Effect of Varenicline and bupropion SR on craving, nicotine withdrawal symptoms, and rewarding effects of smoking during a quit attempt. *Psychopharmacology*, 197:371–377, 2008.
- [16] N.N. Nandola and D.E. Rivera. An improved formulation of Hybrid Model Predictive Control with application to production-inventory systems. *IEEE Transactions on Control Systems Technology*, 21(1):121–135, 2013.