A Novel Engineering Approach to Modeling and Optimizing Smoking Cessation Interventions

by

Kevin Patrick Timms

A Dissertation Presented in Partial Fulfillment
of the Requirement for the Degree
Doctor of Philosophy

Approved November 2014 by the
Graduate Supervisory Committee:

Daniel E. Rivera, Chair
David Frakes
David R. Nielsen

ARIZONA STATE UNIVERSITY

December 2014

ABSTRACT

Cigarette smoking remains a major global public health issue. This is partially due to the chronic and relapsing nature of tobacco use, which contributes to the approximately 90% quit attempt failure rate. The recent rise in mobile technologies has led to an increased ability to frequently measure smoking behaviors and related constructs over time, i.e., obtain intensive longitudinal data (ILD). Dynamical systems modeling and system identification methods from engineering offer a means to leverage ILD in order to better model dynamic smoking behaviors. In this dissertation, two sets of dynamical systems models are estimated using ILD from a smoking cessation clinical trial: one set describes cessation as a craving-mediated process; a second set was reverse-engineered and describes a psychological self-regulation process in which smoking activity regulates craving levels. The estimated expressions suggest that self-regulation more accurately describes cessation behavior change, and that the psychological self-regulator resembles a proportional-with-filter controller. In contrast to current clinical practice, adaptive smoking cessation interventions seek to personalize cessation treatment over time. An intervention of this nature generally reflects a control system with feedback and feedforward components, suggesting its design could benefit from a control systems engineering perspective. An adaptive intervention is designed in this dissertation in the form of a Hybrid Model Predictive Control (HMPC) decision algorithm. This algorithm assigns counseling, bupropion, and nicotine lozenges each day to promote tracking of target smoking and craving levels. Demonstrated through a diverse series of simulations, this HMPC-based intervention can aid a successful cessation attempt. Objective function weights and three-degree-of-freedom tuning parameters can be sensibly selected to achieve intervention performance goals despite strict clinical and operational constraints. Such tuning largely affects the rate at which peak bupropion and lozenge dosages are as-

signed; total post-quit smoking levels, craving offset, and other performance metrics are consequently affected. Overall, the interconnected nature of the smoking and craving controlled variables facilitate the controller's robust decision-making capabilities, even despite the presence of noise or plant-model mismatch. Altogether, this dissertation lays the conceptual and computational groundwork for future efforts to utilize engineering concepts to further study smoking behaviors and to optimize smoking cessation interventions.

DEDICATION

Dedicated to my parents, who instilled within me an appreciation of education and the drive to pursue novel research.

ACKNOWLEDGEMENTS

I would first like to thank Dr. Rivera for all of his guidance since I began my first research rotation in July 2010. He has been an excellent professor, advisor, and mentor. Without his patience, technical expertise, and input regarding training activities, my research would not be of the quality reflected in this document. Furthermore, only with his mentorship was I able to pursue and complete a truly one of a kind, interdisciplinary training experience; his support allowed me to present two posters, give seven presentations at conferences and to collaborators, successfully submit two pre-doctoral fellowship applications, first author three conference papers, two archived abstracts, and three submitted journal manuscripts (to date).

I would also like to thank Dr. David Nielsen and Dr. David Frakes who also served on my dissertation committee. Their input during semesterly committee meetings helped ensure my research and training remained on track and of value.

Furthermore, I owe my gratitude to a number of others who supported me during my graduate school tenure:

Drs. Megan E. Piper at the University of Wisconsin (UW) and Linda M. Collins at Pennsylvania State University (PSU) have been supportive of my training and research activities since I joined the Control Systems Engineering Laboratory. Their letters of support greatly enhanced my fellowship applications. Their input on the technical aspects of this dissertation surely increased the value of this work, and their guidance—as well as that from Dr. Rivera—on how to communicate effectively with diverse disciplinary communities was invaluable.

Dr. JoAnn Williams, Maria Hanlin, and Dr. David Frakes at Arizona State have acted advisors and helped facilitate completion of the degree program. I also appreciate the support of other collaborators at UW and PSU, including Dr. Michael Fiore and Jessica Trail, as well as other Biological Design Graduate Program (BDGP)

leaders and Arizona State professors who helped me successfully pursue a truly unique training experience. Altogether, these individuals have helped me make contributions to engineering, cigarette smoking, and quantitative psychology research literature, and potentially to the care of patients in the future.

I have been fortunate to have had a great number of supporters in the form of colleagues, friends, and family. I would like to thank my colleagues in the Control Systems Engineering Laboratory: Dr. Naresh Nandola, Dr. Sunil Deshpande, Dr. Yuwen (Shirley) Dong, Cesar Martin, Nikhil K.P., and Yunze Yang. I also owe my thanks to a number of other graduate students for listening to my thoughts on research topics and for offering valuable input regarding my training experiences, coursework, research tasks, communicating with diverse research communities, and more. In this regard, I would like to thank Dr. Leah N. Makley, Rene Davis, Dr. Joshua Richer, and Jordan Yaron in particular. Finally, my family has provided endless amounts of encouragement through the years of coursework and research; the support I received from from my mother, Marianne Timms, and sisters, Susan Cantwell, Diane Timms, and Nancy Hein, is immeasurable.

Lastly, I would like to acknowledge the organizations that provided financial support for my training and research: the American Heart Association (pre-doctoral fellowship, 12PRE12030140) and the National Institute on Drug Abuse (NIDA) at the National Institutes of Health (NIH; Ruth L. Kirschstein National Research Service Award Pre-doctoral Fellowship, F31 DA035035); BDGP, the Graduate College, and the Graduate and Professional Student Association at ASU; and the Office of Behavioral and Social Sciences Research and NIDA (R21 DA024266 and K25 DA021173) at the NIH. The work presented in this document does not necessarily represent the views of the organizations that provided funds supporting my training, research, or the resulting presentations and publications.

TABLE OF CONTENTS

CHAPTER Page

CHAPTER Page

CHAPTER Page

LIST OF TABLES

Table Page

LIST OF FIGURES

Figure Page

Figure Page

Figure Page

Figure Page

Chapter 1

INTRODUCTION

1.1 Motivation

Unhealthy behaviors, such as substance abuse, poor diet, and lack of physical activity, remain among the leading causes of preventable death globally (Danaei *et al.*, 2009; Lopez *et al.*, 2006). They also play prominent roles in the development and severity of chronic diseases such as heart disease, chronic obstructive pulmonary disease, diabetes, and more (World Health Organization, 2014; U.S. Department of Health and Human Services and Centers for Disease Control and Prevention, 2014). Through factors including losses in productivity, the cost of prevention programs, and treatment expenses, these deleterious health behaviors have an economic effect on the order of trillions of dollars (World Health Organization, 2005).

Traditionally, preventative or therapeutic interventions intended to promote healthy behaviors are *fixed*. In fixed interventions, all participants receiving the intervention are assigned the same dosage of a treatment component for the entire duration of the intervention (Collins *et al.*, 2004; Rivera *et al.*, 2007). Recently, *adaptive* interventions have emerged as a means to explicitly address the chronic and relapsing nature of many behavioral health disorders, such as substance use dependence (Collins *et al.*, 2004). Part of a growing interest in increasing therapeutic efficiencies through treatment personalization (Fraser *et al.*, 2014; Hamburg and Collins, 2010; Personalized Medicine Coalition, 2014), adaptive interventions generally seek to adjust the components and/or dosages of an intervention over time based on measurements of the patient's circumstances and changing needs (Collins *et al.*, 2004).

Adaptive interventions consist of closed-loop feedback systems, often with feed-forward components (Rivera *et al.*, 2007). This suggests that design of adaptive behavioral interventions can benefit from a control systems engineering perspective. Recent work suggests that well-established controller frameworks such as model-based control and receding-horizon predictive control offer appealing approaches for the design of adaptive behavioral interventions (Deshpande *et al.*, 2014; Dong *et al.*, 2013, 2012; Rivera, 2012; Rivera *et al.*, 2007; Savage *et al.*, in press; Timms *et al.*, 2014d).

Given the gravity of tobacco use as a major global public health issue (Centers for Disease Control and Prevention, 2012; Dube *et al.*, 2010; Erhardt, 2009; Killeen, 2011), this dissertation presents an engineering approach to better understand the process of cigarette smoking behavior change, and proposes an initial framework for an optimized adaptive smoking intervention. Initially, a secondary analysis of data from a smoking cessation clinical trial (McCarthy *et al.*, 2008b) is presented in order to develop dynamic models of the cessation process. The connection between this clinical trial data and engineering modeling methods is made via psychological theory. The conceptual and computational groundwork is then established for a frequently-adapting smoking cessation intervention in the form of a Hybrid Model Predictive control algorithm. This intervention algorithm determines how dosages of counseling and two medications should be varied over time in order to optimally promote a successful cessation attempt for an individual patient.

1.2 Cigarette Smoking: A Persistent Global Public Health Issue

Cigarette smoking is the leading cause of preventable death in the United States, where approximately 440,000 premature deaths and $157B in economic loss are attributed to smoking annually (Dube *et al.*, 2010; Erhardt, 2009; Killeen, 2011). Decades of decreases in U.S. smoking rates have recently stagnated, with approxi-

mately 20% of adults actively smoking (Centers for Disease Control and Prevention, 2012). Globally, smoking is a growing concern and is expected to cause 5 million annual deaths by 2020 (Dube *et al.*, 2010). Cigarette smoking's continued status as a major public health concern is due in part to the chronic and relapsing nature of tobacco dependence (Tobacco Use and Dependence Guideline Panel, 2008; Steinberg *et al.*, 2010; U.S. Department of Health and Human Services and Centers for Disease Control and Prevention, 2014). The fact that over 88% of quit attempts fail (American Cancer Society, 2012) speaks to this assertion.

As quitting smoking is among the single greatest actions an individual can take in order to improve their health (Pipe *et al.*, 2010; U.S. Department of Health and Human Services and Centers for Disease Control and Prevention, 2014), significant resources have been dedicated to designing and evaluating smoking cessation interventions. Such interventions can be behavioral, pharmacological, or community-based (e.g., a public health campaign) in nature (Rivera *et al.*, 2007; Tobacco Use and Dependence Guideline Panel, 2008; U.S. Department of Health and Human Services and Centers for Disease Control and Prevention, 2014). For example, individual counseling that focuses on garnering the ability to cope with cravings or on problem solving skills related to identifying and managing "danger situations" (events or scenarios that may normally lead a person to smoke, e.g., a particular stressor) are particularly effective behavioral therapies (Tobacco Use and Dependence Guideline Panel, 2008). Seven medications have been approved as first line pharmacological tobacco use interventions. Five are nicotine replacement therapies (NRTs), which deliver low doses of nicotine in gum, lozenge, patch, nasal spray, or inhaler form (Rennard *et al.*, 2014; Tobacco Use and Dependence Guideline Panel, 2008). The non-NRT medications are bupropion (Wellbutrin®, Zyban®) and varenicline (Chantix®; Tobacco Use and Dependence Guideline Panel, 2008). Bupropion is an anti-depressant thought to in-

terfere with nicotine addiction mechanisms and reduce craving (Horst and Preskorn, 1998; Warner and Shoaib, 2005). Varenicline is an antagonist for a nicotinic cholinergic receptor and was designed to reduce withdrawal symptoms and minimize the rewarding effect of cigarettes (Benowitz, 2009; West *et al.*, 2008). Evidenced-based clinical practice calls for use of these treatments in combination therapies (except varenicline; Loh *et al.*, 2012; McCarthy *et al.*, 2008b; Piper *et al.*, 2009; Tobacco Use and Dependence Guideline Panel, 2008), such as in a combined bupropion and counseling intervention (Tobacco Use and Dependence Guideline Panel, 2008; McCarthy *et al.*, 2008b; Piper *et al.*, 2009). However, even the most effective combination of pharmacotherapies have 6-month cessation rates below 35% (Tobacco Use and Dependence Guideline Panel, 2008; Piper *et al.*, 2009). Even these modest efficacies are brought into question when considering time periods beyond one year (Irvin *et al.*, 2003; Irvin and Brandon, 2000).

Despite the chronic, relapsing nature of tobacco dependence, current clinical practice generally employs these treatment components in fixed interventions. Furthermore, the U.S. Department of Health and Human services notes that these treatments are used without a well-accepted framework for optimally assigning intervention components to an individual (Tobacco Use and Dependence Guideline Panel, 2008). Altogether, the need for improved smoking cessation interventions is clear.

1.3 Personalizing Behavioral Interventions

The term "personalized medicine" encompasses a broad array of approaches in which treatment is tailored to a specific patient with the general goal of improving treatment outcomes with minimal resource waste. Often, research on personalized interventions focuses on identifying subgroups of patients for whom a specific therapeutic protocol would be most appropriate (Hamburg and Collins, 2010; Personalized

Medicine Coalition, 2014). For example, Bergen *et al.* (2013) examines which medications may be most appropriate for patients with genetic variants related to nicotine receptor sites and neuronal activation.

However, systematic personalization of behavioral interventions remains limited (Collins *et al.*, 2004; Riley *et al.*, 2011; Tobacco Use and Dependence Guideline Panel, 2008). This is due in part to challenges associated with precise identification of causal relationships and quantification of treatment effects. Intuitively, this likely arises from inter-individual variability (Hamburg and Collins, 2010), which may enter at the genetic, biomedical (e.g., co-morbidity), psychological or cognitive, environmental, and socioeconomic levels. The time-varying nature of the cessation process (Chandra *et al.*, 2011; Gwaltney *et al.*, 2005; Rivera, 2012; Timms *et al.*, 2014a,c, 2013; Trail *et al.*, 2014; Velicer and Fava, 2003; Velicer *et al.*, 1992; Walls and Rivera, 2009) makes research on systematic treatment personalization more difficult.

As adaptive interventions attempt to explicitly address the chronic, relapsing nature of disorders, they offer a promising approach for treating behavioral health issues. Such interventions seek to personalize an intervention to a specific patient's circumstances, varying the treatment over time in order to meet the changing needs of individual patients. Specifically, these interventions take the form of *decisions rules*, which systematically assign the components and/or dosages of an intervention based on the levels of the patient's *tailoring variables* (Collins *et al.*, 2004). For example, Kasari *et al.* (2014) considers a two-stage adaptive intervention in which minimally verbal children with autism have a developmental/behavioral intervention that can be augmented at a later point with a speech-generating device, depending on the child's early outcomes.

Recently, the design of adaptive behavioral interventions has been cast as a control systems engineering problem (Bekiroglu *et al.*, 2013; Deshpande *et al.*, 2014; Dong

et al., 2012, 2013, 2014; Nandola and Rivera, 2013; Noble, 2014; Rivera *et al.*, 2007; Savage *et al.*, in press; Timms *et al.*, 2013, 2014d). This approach proposes that decision rules take the form of a control algorithm, with feedback and feedforward signals acting as tailoring variables (Dong *et al.*, 2012, 2013; Riley *et al.*, 2011; Rivera *et al.*, 2007; Timms *et al.*, 2014d). As will be discussed in detail in subsequent sections, employing a control systems paradigm in the design of adaptive behavioral interventions is particularly appealing due to controller functionality related to cohesive handling of multiple manipulated variables and objectives, tunability to influence performance, robustness considerations, and constraint-handling abilities (Morari and Zafiriou, 1989; Nandola and Rivera, 2013; Ogunnaike and Ray, 1994; Rivera *et al.*, 2007). Although ideas related to optimization and control theory have appeared in psychological literature in the past—see Carver and Scheier (1998), Hyland (1987), and Molenaar (1987) for examples—they remain absent from the landscape of current clinical practice when it comes to tobacco use interventions.

1.4 Optimizing Smoking Cessation Interventions: An Engineering Approach

1.4.1 Overview

Although historically applied within traditional engineering settings (Ogunnaike and Ray, 1994), control systems engineering principles have recently been employed within non-traditional contexts including pharmacology (Fraser Health, 2014; Deshpande *et al.*, 2014; Zurakowski and Teel, 2006), environmental engineering (e.g., Brdys *et al.*, 2008; Wang and Garnier, 2012), supply chain management (Nandola and Rivera, 2013; Schwartz and Rivera, 2010), management of anemia (Gaweda *et al.*, 2008), and behavioral health (Dong *et al.*, 2012, 2013; Bekiroglu *et al.*, 2013; Nandola and Rivera, 2013; Noble, 2014; Rivera, 2012; Rivera *et al.*, 2007; Timms *et al.*, 2014d).

Patient care environments have also seen a rise in automation in recent years. Examples of automation ideas being introduced clinically include a clinical trial where a PID controller is used to automate deliver anesthetics (Fraser Health, 2014; Soltesz, 2013), a clinical trial employing fuzzy control for cardiac intensive care patients (Denai *et al.*, 2009; Ross *et al.*, 2009), and a multitude of clinical analyses on artificial pancreas systems drawing from PID and predictive control architectures (Doyle *et al.*, 2014).

Developing an engineering-based smoking cessation intervention draws from a translation of the clinical demands of an intervention into control systems components: treatment goals (e.g., smoking abstinence) correspond to controlled variable set points and targets; tailoring variables correspond to measured outcomes and measured disturbances that act as feedback and feedforward signals, respectively; treatment component dosages correspond to manipulated variables; clinical and pragmatic limits on dosage levels and dosage adjustments correspond to manipulated variable bounds and move size constraints; dose-responses, disturbance-responses, and uncertainties are captured in nominal models; and an individual patient being treated is the "plant" (Rivera *et al.*, 2007; Timms *et al.*, 2014d).

Noted previously, the goal of this research is to lay the computational and conceptual groundwork for an engineering-based adaptive smoking cessation intervention. Toward this goal, two major areas of research are pursued: (1) estimation of dynamical systems models describing the process of behavior change during a quit attempt and (2) formulation, tuning, and analysis of an Hybrid Model Predictive Control (HMPC)-based smoking cessation intervention.

1.4.2 Dynamical Systems Modeling

Behavioral science has a rich history of quantifying the relationship between measured behaviors and related constructs using static modeling methods (Riley *et al.*, 2011), such as structural equation modeling (SEM; Bollen, 1989). This resulted in part from early challenges related to cost-effective and precise measurement of behaviors frequently over time (Collins, 2006; Riley *et al.*, 2011). Largely due to the rise of computerized technologies, behavioral *intensive longitudinal data* (ILD)—measurements of behaviors or related constructs recorded frequently over time, up to continuous-time measurement—are becoming increasingly available (Riley *et al.*, 2011; Timms *et al.*, 2014a; Walls and Schafer, 2006). Mobile and wearable technologies are particularly appealing as platforms for collecting ILD via ecological momentary assessments (EMA), which collect such data in an individual's natural environments (i.e., not clinical or laboratory environments) often in real time (Moskowitz and Young, 2006; Shiffman *et al.*, 2008).

ILD offers an opportunity to estimate dynamical systems models that parsimoniously represent changing behaviors as processes of change in the form of ordinary differential equations (ODEs; Rivera, 2012; Timms *et al.*, 2014a). The estimation of these dynamic models initially relies on well-established system identification methods resulting in rich, comprehensive, and quantitative descriptions of the dynamic characteristics of time-varying behaviors (Ljung, 1999). These nomothetic or idiographic models result in clinically-meaningful parameter estimates, insight into the mechanisms of behavior change, can elucidate novel degrees-of-freedom on which to potentially intervention, can facilitate potentially expanded analyses through simulation, and more (Martín *et al.*, 2014; Rivera, 2012; Timms *et al.*, 2013, 2014a,c). Ultimately, though, modeling cessation with this analytical approach benefits most

significantly from its connection to control systems engineering, where estimated dynamical models can act as the basis for design of control systems-based interventions (Ogunnaike and Ray, 1994; Rivera, 2012; Rivera *et al.*, 2007; Timms *et al.*, 2014a).

As behavioral interventions are often based on hypothesized mechanisms of behavior change, which have historically been described from a static perspective (Riley *et al.*, 2011), recent work has focused on developing dynamical systems models to describe behavior change according to well-known, theorized psychological mechanisms (Dong *et al.*, 2013; Martín *et al.*, 2014; Navarro-Barrientos *et al.*, 2011; Rivera, 2012; Rivera *et al.*, 2007; Timms *et al.*, 2014a,c). To better understand the process of smoking cessation behavior change, dynamical systems models have been developed to describe behavior change thought to occur according to two mechanisms of change: mediation and self-regulation. Specifically, continuous-time dynamical systems models are estimated to examine the role of these mechanisms in cessation drawing from ILD collected in the smoking cessation clinical trial described by McCarthy *et al.* (2008b); Timms *et al.* (2014a,c).

Described in detail in Chapters 2 and 3, the models examined in this dissertation primarily consider *Craving*, patient-reported average daily craving levels, and *CPD*, the total number of cigarettes smoked per day (as recorded nightly from subjects via personal digital assistant [PDA]; McCarthy *et al.*, 2008b). Fig. 1.1 depicts *Craving* and *CPD* ILD for two group averages, one that received active bupropion and counseling ("AC" group) and another that received placebo medication and no counseling ("PNc" group), and for single subject examples from these groups. In Fig. 1.1, the target quit date (TQD) is day 8. Visual inspection of these data makes the case for the utility of an engineering modeling method as *Craving* clearly features inverse response on average and *CPD* features dynamics on two different time scales.

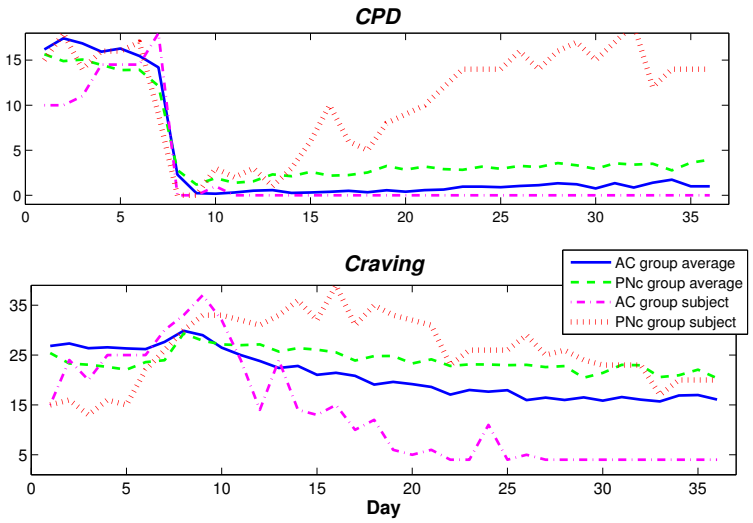

Figure 1.1: *Craving* and *CPD* Signals for the AC Group Average, PNc Group Average, One Subject from the AC Group, and One Subject from the PNc group; ILD Collected in the McCarthy *et al.* (2008b) Study. TQD is Day 8.

Dynamic Mediation Models

Statistical mediation is a central tenant within the social and behavioral sciences. It describes a multivariate relationship in which an independent variable, X, is said to exert its influence on an outcome variable, Y, via changes in a mediator variable, M—which directly affects Y—and a non-mediated path. Statistical mediation is most commonly characterized using static, structural equation modeling (SEM; MacKinnon, 2008; Baron and Kenny, 1986). Navarro-Barrientos *et al.* (2011) illustrates how path diagrams in SEM correspond to steady-state process models of production-inventory systems. Drawing from the connection between SEM and fluid analogies from production-inventory systems (Navarro-Barrientos *et al.*, 2011), a general engineering accounting principle (e.g., conservation of mass) is employed here to describe behavior change as a mediational process (Timms *et al.*, 2014a). Generally, these models describe how M responds over time to changes in X, responding according to some dynamic process represented as P_a; the level of Y responds over time to changes

in the independent variable directly according to some dynamic process represented by $P_{c'}$, as well as through changes in M according to some subprocess represented by P_b. Altogether, from a dynamical systems perspective, mediation consists of a parallel cascade process structure (Timms *et al.*, 2014a).

An initial set of dynamic mediation models were estimated that consider initiation of a quit attempt (represented by *Quit*, a unit step on TQD) as X, *Craving* as M, and *CPD* as Y (Timms *et al.*, 2014a). Employing a prediction-error approach (Ljung, 1999), continuous-time transfer functions $P_a(s)$, $P_b(s)$, and $P_{c'}(s)$ were estimated for each treatment group average and the two single subjects. The group average models accounted for up to 87.77% of the variance in the *Craving* signal, and up to 91.49% of the variance in the *CPD* signal. The structures and parameters of these estimated models reflect the major dynamic features observed in each group, such as the inverse response in *Craving*, dramatic reduction in *CPD* on TQD, and slow and small resumption of smoking after TQD. Comparing the parameter estimates across the groups suggests both bupropion and counseling affect cessation dynamics. For example, the parameter estimates indicate the relative degree to which *Craving* increases on the TQD before ultimately decreasing is smallest for the AC treatment condition and greatest for the PNc treatment condition.

A second set of mediation models were estimated using the group average data in which *Quit* was the independent variable, *CPD* the mediator, and *Craving* the outcome. These estimated models resulted similarly high goodness-of-fit values. This suggests that the relationship between *Craving* and *CPD* cannot be fully explained by statistical mediation alone (Timms *et al.*, 2014a).

Dynamic Self-Regulation Models

The concept of self-regulation has historically been of interest as a potential mechanism of behavior change (Carver and Scheier, 1998; Solomon, 1977; Solomon and Corbit, 1974; Velicer *et al.*, 1992; Walls and Rivera, 2009), although data collection challenges in the pre-personal computing era made preceise modeling of this psychological phenomenon difficult (Collins, 2006; Riley *et al.*, 2011). *Craving* and *CPD* ILD affords the opportunity to examine a self-regulatory relationship between these constructs (Collins, 2006; Rivera, 2012; Timms *et al.*, 2014a,c; Walls and Rivera, 2009). Furthermore, analysis of such a relationship significantly benefits from a control systems engineering perspective, given this field's emphasis on understanding and manipulating closed-loop processes (Ahmed and Koob, 1998; Dong *et al.*, 2012, 2013; Ogunnaike and Ray, 1994; Timms *et al.*, 2014a,c).

Closed-loop models were developed describing cessation as the process depicted in Fig. 1.2. In this mechanism, *Craving* acts as the primary outcome, the level of which is directly a result of smoking, i.e., the level of *CPD*. Changes in *CPD* are the result of a disturbance path and a feedback path: *Quit* is the disturbance that acts as an independent, exogenous influence on *CPD*; changes in *CPD* due to the latter path result from a psychological or biological self-regulator, which promotes or deters smoking in order to track a *Craving* set point (r_{crav}; Timms *et al.*, 2014a,c). For each set of group average ILD from the McCarthy *et al.* (2008b) clinical trial, three continuous-time ODEs were estimated using a prediction-error approach (Ljung, 1999, 2011; Timms *et al.*, 2014a). The closed-loop identification problem consists of a two step estimation procedure (as is also required when estimating dynamic mediation models): one transfer function, $P(s)$, is estimated as part of a single-input single-output (SISO) problem in which *CPD* is the input signal and *Craving* the output;

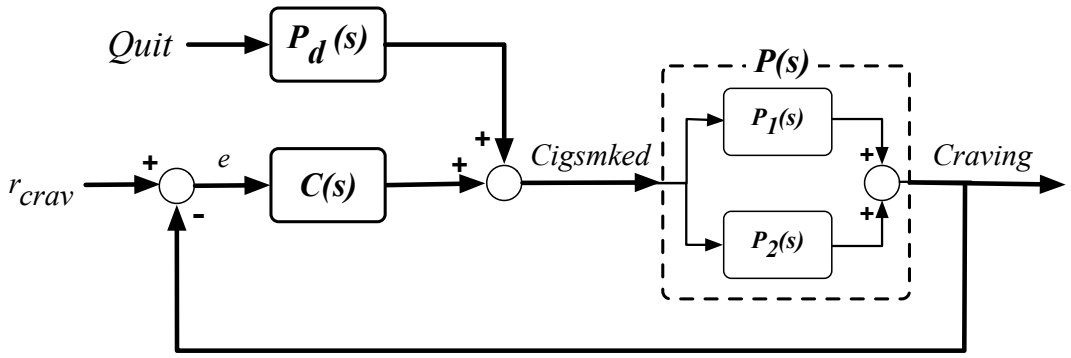

Figure 1.2: Block Diagram Describing the Process of Smoking Cessation Behavior Change as a Self-Regulated System (Timms *et al.*, 2014a,c).

two transfer-functions, $P_d(s)$ and $C(s)$, are simultaneously estimated as part of a two-input one-output problem where *Quit* and the difference between *Craving* and r_{crav} are the inputs, respectively, with *CPD* as the output. The appropriate r_{crav} value was empirically determined to be the average pre-TQD *Craving* level (Timms *et al.*, 2014a).

As with the mediation models, high goodness-of-fit values were obtained for the group averages—up to 87.44% fit for *Craving* and 91.44% for *CPD*—which are obtained with low-order transfer function structures. Notably, $P(s)$ is adequately represented with a first order with zero transfer function. Further analysis indicates this function accounts for sum of two parallel subprocesses; one corresponds to the immediate increase observed in *Craving* on TQD and the second is an opposing subprocess that brings *Craving* down to net lower values as time goes on. The $P_d(s)$ function captures the drop in *CPD* on TQD. $C(s)$ was found to be represented by a first order continuous-time transfer function (Timms *et al.*, 2014a). This suggests that on average, the self-regulator behaves as a proportional-with-filter controller (Morari and Zafiriou, 1989; Timms *et al.*, 2014a, 2013).

1.4.3 Control Systems Engineering for Smoking Interventions

The smoking cessation decision framework detailed in the following involves intensive measurement and intensive adaptation—both primarily on a daily basis. Consequently, this constitutes a highly personalized intervention. Fig. 1.3 visually depicts the general location such an intervention takes in the "intervention optimization space." As seen in Fig. 1.3, this type of intervention contrasts a typical nicotine patch intervention protocol; patch dosage is typically determined by whether a patient smokes at least 10 or 15 cigarettes per day, and that daily patch dosage is generally maintained from TQD throughout the duration of the intervention (Tobacco Use and Dependence Guideline Panel, 2008; Rennard *et al.*, 2014). It should be noted that each area in this space features various opportunities, advantages, and clinical challenges.

This dissertation proposes an adaptive smoking intervention algorithm that systematically determines daily adjustments to treatment components based on known intervention targets, feedback and feedforward signals, and an understanding of the cessation process and clinical constraints. Hybrid Model Predictive Control (HMPC) offers a distinctly appropriate framework for such an intervention:

- Control action is determined by minimizing a quantified optimality criterion (Camacho and Bordons, 1999; Goodwin *et al.*, 2005; Nandola and Rivera, 2013; Rossiter, 2003), which is conductive to an *optimized* personalized intervention (Nandola and Rivera, 2013; Rivera *et al.*, 2007; Timms *et al.*, 2014d).

- HMPC considers a dynamical system with logical and/or discrete components and constraints (Bemporad and Morari, 1999). Optimal dosing assignments can therefore be determined despite the fact that the manipulated variables, i.e., the treatment components, can only assume pre-determined, discrete levels

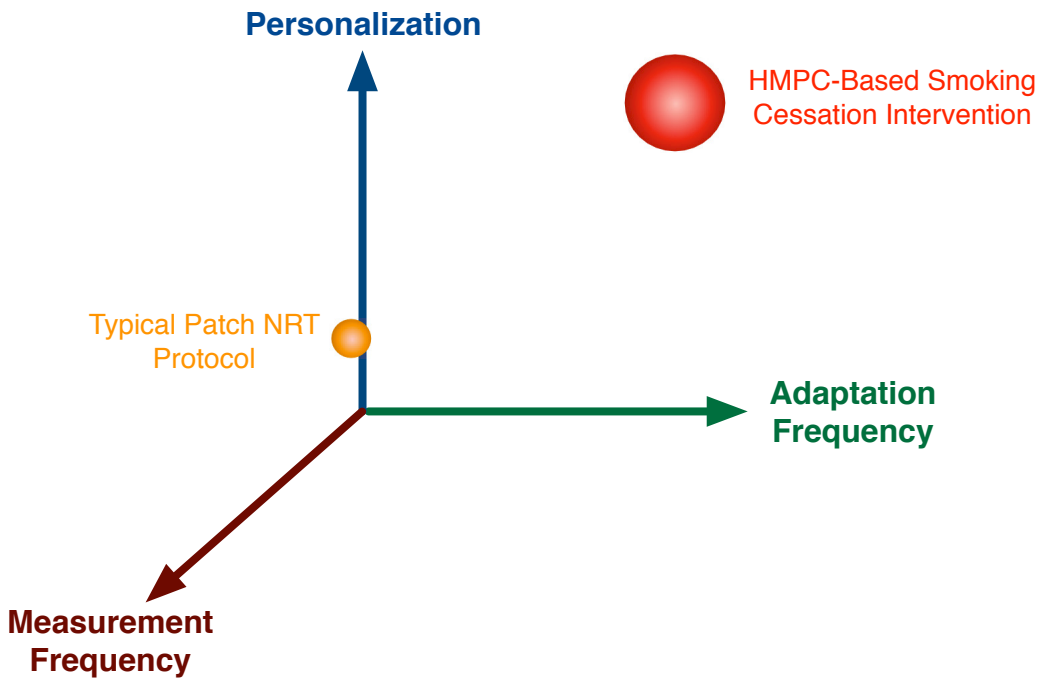

Figure 1.3: Conceptual Depiction of the Intervention Optimization Space, Including the Approximate Locations of a Typical Nicotine Patch Protocol and the HMPC-Based Adaptive Intervention Formulated in this Document. The Origin Corresponds to an Intervention that is Not Personalized to Any Degree, Involves no Formal Measurement, and is Not Varied Over Time.

(Deshpande *et al.*, 2012; Dong *et al.*, 2013; Nandola and Rivera, 2013; Timms *et al.*, 2014d). This means that an HMPC-based intervention algorithm can only assign bupropion in 150 mg doses, for example (Lexicomp, 2014; Timms *et al.*, 2014d; Tobacco Use and Dependence Guideline Panel, 2008).

- Combined feedback and feedforward action means dosage assignment is based on measurements of a patient's changing needs (Nandola and Rivera, 2013; Rivera *et al.*, 2007; Timms *et al.*, 2014d), which is the fundamental motivation for use of adaptive interventions in behavioral health (Collins *et al.*, 2004).

- The feedback and feedforward features help facilitate effective dosing despite unmeasured disturbances and potentially significant patient-to-patient variabil-

ity (i.e., plant-model mismatch; Camacho and Bordons, 1999; Goodwin *et al.*, 2005; Nandola and Rivera, 2013; Rossiter, 2003). In other words, the ability to tune for robust performance supports the clinical-relevance of such an engineering-based intervention.

- The ability to tune an HMPC formulation—particularly one with three-degree-of-freedom features—can allow a clinician to flexibly balance outcome goals with dosing demands (Dong *et al.*, 2013; Deshpande *et al.*, 2014; Nandola and Rivera, 2013).

- The feedback, feedforward, and predictive nature of HMPC (and Model Predictive Control [MPC] more generally; Camacho and Bordons, 1999; Goodwin *et al.*, 2005; Nandola and Rivera, 2013; Rossiter, 2003) means dosing decisions are made with a patient's past, present, and future needs in mind (Deshpande *et al.*, 2014; Dong *et al.*, 2013; Nandola and Rivera, 2013; Rivera *et al.*, 2007; Timms *et al.*, 2014d). Therefore, an intervention of this nature is conducive to "just-in-time" decision making (Timms *et al.*, 2014d), which means that dosage adjustments can be assigned in order to avoid future conditions that may lead to deviations from treatment targets (Intille *et al.*, 2003; Kumar *et al.*, 2013; Riley *et al.*, 2011).

- HMPC is well-suited for multi-input multi-output (MIMO) systems (Camacho and Bordons, 1999; Goodwin *et al.*, 2005; Nandola and Rivera, 2013; Rossiter, 2003), which is conducive to systematic adaptation of interventions featuring multiple treatment components and multiple objectives. This is critical, given the push for combination therapies (Tobacco Use and Dependence Guideline Panel, 2008) and the fact that a variety of risk factors are thought to promote smoking relapse (Cui *et al.*, 2006; Shiffman, 2005; Timms *et al.*, 2014d; Tobacco

Use and Dependence Guideline Panel, 2008; Yasin *et al.*, 2012).

- HMPC explicitly considers hard constraints while determining manipulated variable adjustments (Camacho and Bordons, 1999; Goodwin *et al.*, 2005; Nandola and Rivera, 2013; Rossiter, 2003) implying dosing decisions can adhere to clinically-necessitated restrictions (e.g., medication toxicity levels) as well as implementation considerations (e.g., resource limitations determined by a patient's health insurance policy).

- Because solutions to the optimization problem within HMPC computations involve well-established routines (Goodwin *et al.*, 2005; Holmstrom *et al.*, 2009; The MathWorks, 2014d; Nandola and Rivera, 2013; Rossiter, 2003), calculation of dosing decisions is computationally tractable and can be implemented within existing technological infrastructure (potentially even a mobile phone). These solutions take the form of a quadratic program (QP) or mixed integer quadratic program (MIQP) (Holmstrom *et al.*, 2009; Nandola and Rivera, 2013).

Altogether, these factors suggest that design of an frequently-adapting smoking cessation intervention from intensively sampled data can be pursued as a controller design problem.

Primary Intervention Architecture

The controller design problem presented in this dissertation primarily considers the structure depicted in Fig. 1.4. The controlled variables are *CPD* and *Craving*, considered to be self-reported via a smart phone application. The target levels of both outcomes are 0 as of TQD. Consequently, *CPD* and *Craving* act as feedback signals, while *Quit* and *Stress* (average daily patient-reported stress level) are feedforward

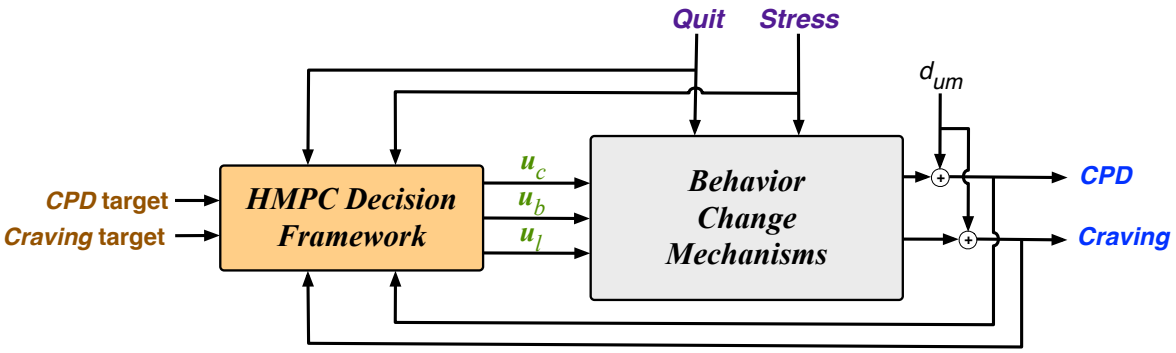

Figure 1.4: Block Diagram of the General Decision Framework for an Adaptive, Smoking Cessation Intervention that Employs an HMPC Algorithmic Structure to Determine Daily Dosages of Counseling (u_c), Bupropion (u_b), and Nicotine Lozenges (u_l). *CPD* and *Craving* are Controlled Variables, *Quit* and *Stress* are Measured and Anticipated Disturbance Variables, and d_{um} Represents an Unmeasured Disturbance.

signals that are measured and anticipated. The intervention algorithm assigns daily dosages of counseling, bupropion, and nicotine lozenges.

As described in Chapter 3, the intervention formulation involves the following. The nominal models that act as the basis for the controller reflect the open-loop dynamics for a representative yet hypothetical patient who unable to quit smoking on their own. The open-loop *CPD* and *Craving* responses to initiation of a quit attempt are patterned after dynamic self-regulation models estimated for the previously-alluded-to subject from the PNc group in the McCarthy *et al.* (2008b) study. These models reflect the fact that the *Quit* disturbance initially moves *CPD* toward the intervention goal and simultaneously moves *Craving* further from its target, as depicted with the dotted red data in Fig. 1.1. These signals reflect an intuitive *CPD*-*Craving* interrelationship. The patient is initially able to reduce smoking just after the quit attempt, but doing so corresponds to an increase in *Craving*. However, this patient is unable to sustain the initial abstinence, with *CPD* gradually increasing to approximately pre-TQD levels as *Craving* simultaneously settles to pre-TQD levels (Timms *et al.*, 2014a). The open-loop dose-response and *Stress*-response models are

informed by data, literature, and simulation.

As the manipulated variables can only be assigned in pre-determined, discrete-valued levels, the behavior change dynamical process and constraints are represented as a mixed logical dynamical structure (Bemporad and Morari, 1999; Nandola and Rivera, 2013). Continuous and discrete auxiliary variables translate the descrete dosage levels into linear inequality constraints (Bemporad and Morari, 1999), enforcing that counseling can only be assigned as 0 or 1 sessions per day, bupropion as 0, 1, or 2 150 mg doses per day (Lexicomp, 2014; Tobacco Use and Dependence Guideline Panel, 2008), and integer values between 0 and 20 lozenges per day (Rennard *et al.*, 2014). Daily decision-making is done in the context of a number of constraints in order to adhere to clinical requirements and practical considerations. For example, a bupropion move size constraint is defined to allow only unit increases in bupropion dosage at a time, reflective of the on-label use of bupropion and current clinical recommendations (Lexicomp, 2014; Tobacco Use and Dependence Guideline Panel, 2008).

As TQD is typically pre-determined (Tobacco Use and Dependence Guideline Panel, 2008; Lexicomp, 2014), the set points changes and *Quit* disturbance are anticipated during each decision point. With a multi-week prediction horizon, these changes are anticipated from the first decision period. *Stress* is also anticipated in this formulation, where *Stress* is assumed to take on the most recent measured value for the duration of the prediction horizon. Here, *Stress* is represented as an stochastic, auto-correlated signal (DeLongis *et al.*, 1988; Shiffman and Waters, 2004).

As a simple illustration of the detailed analyses in Chapter 4, Fig. 1.5 depicts the disturbance signals, dosing decisions as determined by the controller, and the corresponding *CPD* and *Craving* responses (predicted through simulation) when the patient receiving the intervention is that around whom the intervention was designed

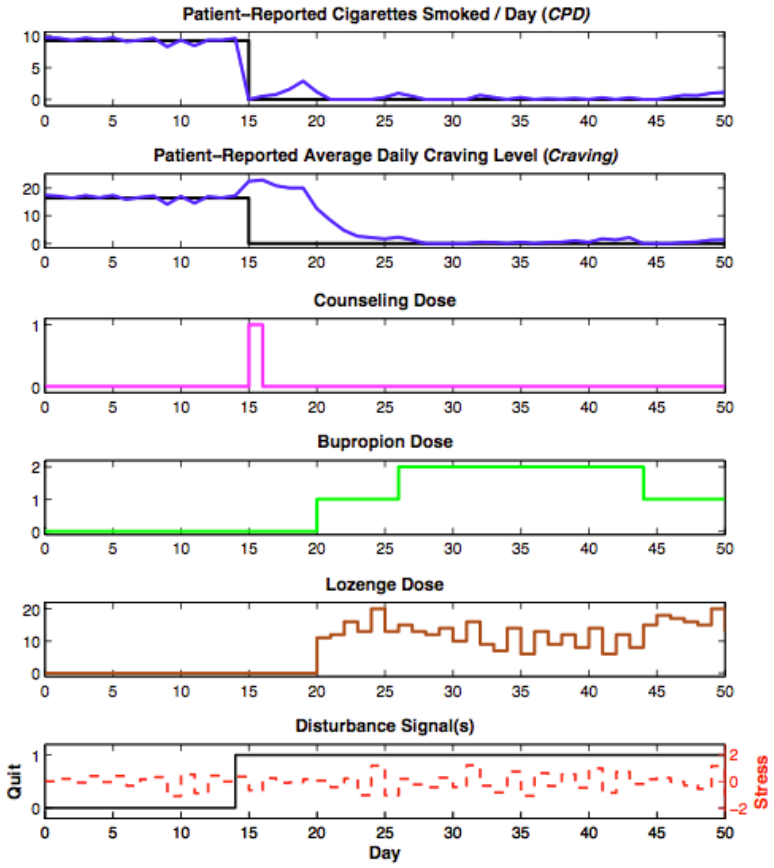

Figure 1.5: *Craving* and *CPD* Responses to *Quit* and the Corresponding Set Point Changes (Solid Black), the Stochastic *Stress* Disturbance (Dashed Red), and the Depicted Counseling (Magenta), Bupropion (Green), and Lozenge (Brown) Dosing, which is Determined by the HMPC-Based Intervention. (The Objective Function Penalty Weight for the *CPD* Target Equals 10, and that for the *Craving* Target Equals 1.)

(i.e., negligible to no plant-model mismatch). This scenario considers two weeks pre-TQD (i.e., TQD is day 15) through the first five weeks of the quit attempt. Here, the controller employs an 8 day move horizon and a 30 day prediction horizon, and the objective function penalty weight for *CPD* target tracking is an order of magnitude larger than that for the *Craving* target.

Fig. 1.5 indicates that under these conditions, both intervention targets are ultimately tracked to a successful degree. Dosing is initially more hesitant than most

current treatment protocols, with minimal dosing around TQD. This reflects the fact that the patient is briefly able to abstain to a relatively successful degree on their own. The aggressive bupropion and lozenge dosing implemented within the first week of the quit attempt serves to suppress the full relapse that would otherwise occur, and to bring down *Craving* to approximately target levels within 10 days of TQD. Visual inspection of Fig. 1.5 indicates that the lozenge treatment, which has an immediate effect, as does *Stress* (both sets of open-loop transfer function models are semi-proper), has high variance after day 20 in order to reject the stochastic *Stress* disturbance. Overall, some lapses may occur, but the quit attempt remains relatively successful: less than 15 cigarettes are smoked in the five weeks after TQD, nearly three of those weeks are smoke-free (non-consecutively), and the maximum lapse is less than three cigarettes.

Fig. 1.5 corresponds to a relatively straightforward formulation and a minimally-tuned scenario. Chapter 3 incorporates and evaluates the effect of additional complexity to enhance the formulation in a clinically-meaningful manner. For example, the objective function penalizes offset for a third set point in order mitigate overly-aggressive dosing demands. Three-degree-of-freedom tuning functionality is also integrated into the HMPC architecture which facilitates more intuitive tuning. Analysis of nominal performance of the intervention for various tuning scenarios first demonstrates the flexibility of this treatment strategy for obtaining dosing that promotes smoking abstinence and feature various treatment schedule characteristics; as will be shown, the inter-related nature of *Craving* and *CPD* supports relatively good and consistent abstinence outcomes even with significant detuning meant to influence the trade off between daily dosing peaks and treatment adjustment intensity. Robust performance is also evaluated when *CPD* and *Craving* feature measurement noise and plant-model mismatch in the form of both parametric uncertainty and non-parametric

uncertainty in which the patient receiving the intervention is modeled after a subject in the McCarthy *et al.* (2008b) study who is not the nominal patient.

1.5 Contributions of This Research

This dissertation documents two closely-related areas of research: (1) use of dynamical systems modeling and system identification techniques to better understand behavior change during smoking cessation; (2) development and evaluation of an HMPC-based decision system for an optimized, adaptive smoking intervention. Some of the model development and identification research entails the following.

- Development of dynamical systems representations of psychological theories. Dynamical systems models describing hypothesized mechanisms of change that are prominent within social science, behavioral science, and substance use literature—statistical mediation and self-regulation—are developed that can offer novel insight into these hypothesized mechanisms.

- An illustration of what linear, continuous-time dynamical systems models and system identification methods can offer in terms of better understanding health behaviors, particularly tobacco-use behaviors.

- Application of an engineering approach to describe statistical mediation as a time-varying process of behavior change, which can be examined in the context of a traditional, cross-sectional behavioral science experiment. Secondary analysis of clinical trial data (McCarthy *et al.*, 2008b) is used in a case study of dynamic mediation modeling.

- The use of a control systems engineering perspective to understand the role of self-regulation within day-to-day changes in smoking and craving levels during an attempt to quit smoking. In secondary analysis of the McCarthy *et al.*

(2008b) study, a closed-loop identification problem is pursued, resulting in estimated models that benefit from a PID-control interpretation.

- Brief demonstration of how these engineering modeling techniques can be used to provide insight into the manner in which mediated and self-regulatory behavior change is affected by treatment conditions (through nomothetic model estimation), as well as to provide insight into intra-individual effects (through idiographic model estimation).

The intervention-design aspect of this research consists of a controller design problem that acts as a proof-of-concept for the potential clinical utility of an engineering-based adaptive smoking intervention. The nature of the smoking cessation problem necessitates an engineering-based treatment strategy with a unique combination of demands and features compared to similar controller design efforts in alternate behavioral health interventions. For example, the cessation intervention revolves around a central event—initiation of the quit attempt. The TQD can be anticipated weeks in advance, and the quit attempt acts as a measured and anticipated step disturbance which directly corresponds to step changes in the intervention targets. This disturbance initially moves *CPD* toward its target, and the *CPD* target distinctly switches from smoking "on" to smoking fully "off." Because of such considerations, major components of this research include the following.

- A set of requirements for an HMPC-based smoking intervention is determined such that the intervention proposed here closely matches clinical reality, while managing pragmatic challenges associated with clinical implementation.

- A clinically-relevant adaptive intervention architecture (e.g., Fig. 1.4) is proposed that employs a control systems paradigm to determine day-to-day adjust-

ments to treatment component dosages in order to optimally meet the changing needs of a patient.

- An intervention algorithm is formulated employing an HMPC structure. This formulation translates the requirements of an intensive adaptive cessation intervention (alluded to previously) into a state-space representation of the system that adheres to a mixed logical dynamical framework, linear inequality constraints, and an objective function. Computation of control action takes the form of a MIQP (Holmstrom *et al.*, 2009; Nandola and Rivera, 2013).

- The closed-loop intervention is evaluated through simulation. A tuning strategy is developed by analyzing the nominal and robust performance of the intervention algorithm in both the absence and presence of three-degree-of-freedom (3DoF) tuning capabilities. Ultimately, the case is made for the potential clinical utility of this intervention approach, which features significant flexibility in order to obtain various clinically-meaningful dosing demands and controlled variable performance.

Details on future efforts that could move this approach toward clinical implementation are also provided, and the overall findings from this research are summarized.

1.6 Outline

This chapter provided an introduction to the problem setting and motivation for this research. The approach, methods, and findings presented in detail in the following chapters were also briefly outlined.

Chapter 2 further motivates and generally describes the utility of dynamical systems modeling and system identification methods for understanding health behaviors. The nature of ILD and its role in understanding smoking behavior change are out-

lined. Details on a continuous-time system identification approach for estimating and interpreting mediation and self-regulation models are presented; estimation draws from ILD collected nightly in the McCarthy *et al.* (2008b) study. This model development effort is presented here largely through incorporation of the Timms *et al.* (2014a) manuscript, which cohesively summarizes the modeling and estimation methods and findings. Extensions of the concepts and opportunities presented in Timms *et al.* (2014a) are also provided, including an illustration of opportunities for exploring therapeutic mechanisms through simulation and investigation of the role of self-regulation on a within-day level.

Chapter 3 formulates the HMPC-based intervention. First, connections between the clinical demands of a cessation intervention and control systems engineering ideas are established. Details on the nominal patient models, hybrid linear models of the overall system, control action computations, and the 3DoF structure are presented.

Chapter 4 evaluates the performance of the proposed HMPC-based intervention through diverse simulations. Nominal performance is first considered. For this, the effects of adjusting the objective function penalty weights and 3DoF tuning parameters on the character of dosing demands and post-TQD $Craving$ and CPD responses are evaluated. Robust performance is then considered when noise and plant-model mismatch in the form of parametric uncertainty and alternate patient models from the McCarthy *et al.* (2008b) study are introduced.

Chapter 5 summarizes summarizes the contribution and overall findings of this research.

1.6.1 Publications

Much of the material presented in this dissertation has also been incorporated into a series of publications. The following primarily describe modeling aspects of

this research:

- K.P. Timms, D.E. Rivera, L.M. Collins, and M.E. Piper, "System identification modeling of a smoking cessation intervention", Proceedings of the 16th IFAC Symposium on System Identification pp. 786-791 (2012).

- K.P. Timms, D.E. Rivera, L.M. Collins, and M.E. Piper, "Control systems engineering for understanding and optimizing smoking cessation interventions", Proceedings of the 2013 American Control Conference pp. 1967-1972 (2013).

- K.P. Timms, D.E. Rivera, L.M. Collins, and M.E. Piper, "Continuous-time system identification of a smoking cessation intervention", International Journal of Control **87**, 7, 1423-1437 (2014).

- K.P. Timms, D.E. Rivera, L.M. Collins, and M.E. Piper, "A dynamical systems approach to understanding self-regulation in smoking cessation behavior change", Nicotine and Tobacco Research **16**, Suppl. 2, S159-S168 (2014).

- K.P. Timms, C.A. Martín, D.E. Rivera, E.B. Hekler, and W. Riley, "Leveraging intensive longitudinal data to better understand health behaviors", Proceedings of the 36th Annual IEEE EMBS Conference pp. 6888-6891 (2014).

- K.P. Timms, D.E. Rivera, L.M. Collins, and M.E. Piper, "Dynamic modeling and system identification of mediated behavior change with a smoking cessation case study", Multivariate Behavioral Research (*Submitted*).

An early version of the HMPC-based controller is briefly developed and evaluated in:

- K.P. Timms, D.E. Rivera, M.E. Piper, and L.M. Collins, "A Hybrid Model Predictive Control strategy for optimizing a smoking cessation intervention", Proceedings of the 2014 American Control Conference pp. 2389-2394 (2014).

Most of the analyses in Chapter 4 have not yet been incorporated into manuscript form; anticipated publication venues for this work include Control Engineering Practice and the Journal of Consulting and Clinical Psychology.

Chapter 2

DYNAMICAL SYSTEMS MODELING FOR BETTER UNDERSTANDING TIME-VARYING HEALTH BEHAVIORS

2.1 Overview

Historically, the study of dynamic behavior relied on experiments in tightly-controlled laboratory environments, epidemiological correlation analyses, and randomized, cross-sectional (e.g., factorial), placebo-controlled clinical trials (Hekler *et al.*, 2013a; Shadish *et al.*, 2002). The prominence of these experimental approaches was due in part to challenges with cost-effective data collection in non-controlled settings and frequently over time. Relatedly, methodological and analytical efforts focused on research questions that could be answered with relatively few, non-repeating observations (Walls and Schafer, 2006).

The personal computing era has led to a rapid rise in the availability of intensive longitudinal data (ILD; Collins, 2006; Riley *et al.*, 2011; Walls and Schafer, 2006). The increased ability to collect data sets with such intensive measurements has largely outpaced development of analytical methods optimally suited to characterize the extensive phenomena captured within these data (Walls and Schafer, 2006). The intensive nature of these data offer an opportunity for introduction engineering modeling methods into social and behavioral science research settings. This chapter illustrates how valuable insight into the dynamic nature of health behaviors can be elucidated using dynamical systems modeling and system identification methods, which are commonly used within engineering settings to model dynamic processes.

2.2 ILD: Collection & Data Types

The term *intensive longitudinal data* is used to describe data sets where observations are "recorded at more than a handful of time points," up to continuous-time measurement of events. The nature of these intensive measurements can vary widely: ILD may be quantitative or qualitative in nature, measured at consistent or irregular intervals, involve fixed or random measurement, and more. Simple examples of ILD include number of steps per day as recorded by a pedometer and surveys of emotional states completed by individuals within a single day (Walls and Schafer, 2006)

Generally, ILD offers an opportunity to study dynamic behavioral phenomena (Collins, 2006; Walls and Schafer, 2006). In doing so, one of the greatest opportunities afforded by ILD is the study of *intra*-individual variability (Collins, 2006; Molenaar, 2004). This is generally a departure from traditional behavioral science analyses, which has historically examined cross-sectional intervention data sets primarily (Collins, 2006). More generally, ILD facilitates a time-varying analytical perspective when comparing groups. For example, Timms *et al.* (2014c) develops separate models of a single dynamic process as it is observed in two different treatment groups. Also, ILD have largely facilitated the rise of a research paradigm focusing on single subjects (Molenaar and Campbell, 2009; Molenaar, 2004). Altogether, the rise in the use and availability of ILD has opened up a number of opportunities for more comprehensive modeling and understanding of health behaviors.

Ecological momentary assessment (EMA) encompasses a range of experimental approaches in which ILD is collected in the context of an individual's natural environment—i.e., not a laboratory or clinical environment (Shiffman *et al.*, 2008). For example, patient self-reports of withdrawal symptoms as collected via a smartphone application from subjects trying to quit smoking constitute EMA (McCarthy

et al., 2008b, 2006; Shiffman *et al.*, 2008). EMA now plays a larger role in the study of behavior due in part to the recent rise in advanced mobile and wearable technologies (Hekler *et al.*, 2013b). Furthermore, the appeal of ecological momentary interventions with "just-in-time" decision-making motivates greater focus on EMA: measurement of behaviors and related constructs in a patient's natural environment can directly inform adaptation of an intervention such that it can most effectively address a patient's ecological and situational circumstances at that moment (Heron and Smyth, 2010).

2.3 Dynamic Modeling & System Identification for Studying Health Behaviors

2.3.1 Overview

The rate at which methods for modeling ILD have been developed has lagged behind the rate at which such data has become available (Walls and Schafer, 2006). In a departure from a long tradition of static modeling methods frequently employed by quantitative psychologists (e.g., structural equation modeling; Bollen, 1989), attention is increasingly given to development of modeling methods that extract the additional information captured in ILD. Many of which are based in systems science ideas, some of these emerging approaches are described in Boker (2012), Boker and Nesselroade (2002), Ginexi *et al.* (2014), Riley *et al.* (2011), Rivera (2012), Tan *et al.* (2012), Trail *et al.* (2014), and Timms *et al.* (2014a).

Dynamical systems modeling and system identification methods from engineering offer a means to more comprehensively characterize the behavioral dynamics at play during an attempt to quit smoking (Timms *et al.*, 2014a,c). In fact, the maturity and flexibility of these principles—and ultimately their connection to the design of closed-loop interventions—has resulted in their recent use in answering important

questions outside of traditional engineering settings. Examples include modeling dynamic cellular processes (Ahn *et al.*, 2006), the effects of rainfall (Garnier and Young, 2014), supply chains (Schwartz and Rivera, 2010), climate change (McGuffie and Henderson-Sellers, 2005), economics (Franksen, 1969), and crowd control (Schwager, 2008).

Incorporating engineering ideas and perspectives into the study of psychological phenomena is not limited to recent work (Carver and Scheier, 1998; Hyland, 1987; Molenaar, 1987). For example, McCleary and Hay (1980) and McDowall *et al.* (1980) describe how time-series models can represent dynamic phenomena observed in data from an "interrupted time series;" an interrupted time series consists of sampled data from before and after introduction of an intervention or stimuli that has induced a change in the dependent variable; this type of study is considered to be a "quasi-experiment" in behavioral science settings (Shadish *et al.*, 2002). Specifically, McCleary and Hay (1980) discusses an "ARX with intervention" model in which an intervention is represented as an exogenous input; if the coefficient(s) for these intervention terms are statistically significant, the case is made for the causal effect of the intervention on the dependent variable dynamics. A more recent effort documented in Khuder *et al.* (2007) employs an ARX with intervention model to characterize how a clean air ordinance in Bowling Green, Ohio, affected local rates of coronary heart disease.

2.3.2 Dynamical Systems Modeling

Dynamical systems models mathematically represent processes of change. Specifically, they represent how an output variable in a system (y) responds over time to changes in an input variable (u) or a disturbance variable (d); u is typically known and controlled, while d can be a measured or unmeasured exogenous factor (Ljung,

1999; Ogunnaike and Ray, 1994). In behavioral science contexts, u correspond to independent or predictor variables and y are dependent variables that correspond to outcomes of interest (Timms *et al.*, 2014c).

Dynamical systems modeling encompasses a broad array mathematical representations; these expressions may be continuous or discrete, lumped or distributed parameter, deterministic or stochastic, linear or nonlinear, and so on (Ljung, 1999). The different general characteristics of the models determine the form of these mathematical functions. For example, lumped parameter representations "lump" all process dynamics into a single variable, t, and are represented by ordinary differential equations (ODEs). Distributed parameter models account for spatial and temporal dynamics separately and are represented by partial differential equations (Ogunnaike and Ray, 1994). Continuous-time models are written in the time or frequency domains; models of the latter case are often dealt with in transfer function form. Alternatively, discrete-time models often take the form of transfer functions in the z domain, as obtained via the $\mathcal{Z}$-transform, i.e., the discrete-time equivalent of the Laplace-transform; these models are also frequently written in the time-domain as a sampled, difference equation in which the output at time k, $y(k)$, is written as a function of y at previous time instances ($y(k-1)$, ..., $y(k-n_y)$), and of the value of the input at current or previous time instances ($u(k)$, ..., $u(k-n_u)$; Franklin *et al.*, 1998; Ljung, 1999).

The general approach to modeling can also vary. Black-box models represent input-output dynamics in a way that is agnostic of the exact mechanisms by which changes in input variables lead to changes in output variables. Models derived from first principles describe processes according to the physical laws in which the process is based. Of these two approaches, the former is a purely data-driven approach, while the latter is theory-driven (Ljung and Glad, 1994).

Frequently, dynamic models fall within the extreme bounds of fully black-box or

first principles. A common starting point for dynamic model development relies on a general accounting principle:

$$\text{Accumulation} = \text{In} - \text{Out} \tag{2.1}$$

In physical systems such as unit operations in engineering settings, equation 2.1 generally corresponds to conservation of mass and energy equations (Felder and Rousseau, 2004; Ljung, 1999; Ogunnaike and Ray, 1994). However, this general principle has also guided model development in settings related to supply chain management (Schwartz and Rivera, 2010; Nandola and Rivera, 2013), pain management (Deshpande *et al.*, 2014; dos Santos *et al.*, 2013), physiology (Dong *et al.*, 2013; Hall *et al.*, 2012; Hekler *et al.*, 2013a), economics (Franksen, 1969), climate change (McGuffie and Henderson-Sellers, 2005), social science (Ionides *et al.*, 2006), crowd control (Schwager, 2008) and more.

Ultimately, choice of the model form and modeling approach are determined by the needs of the modeler. This decision frequently depends on data availability, tractability of model estimation and management, and the planned end use of the model. From a control systems engineering perspective, the modeling task should identify the simplest model of the process that is "good enough" to inform effective control action (Rossiter, 2003).

2.3.3 Dynamic Models for a Smoking Cessation Intervention

This dissertation primarily models the cessation process using ILD from the McCarthy *et al.* (2008b) clinical trial. This was briefly introduced in Chapter 1 and is further described in this chapter. More specifically, ILD was collected through EMA using personal digital assistants, in which participants self-reported withdrawal symptoms, cravings, and affect variables on 10-point Likert scales: 1-11, *No!! ... Yes!!*

(McCarthy *et al.*, 2008b). Intuitively, self-reported data is often noisy and subject to potential data quality issues, such as recall bias or cognitive and situational influences (Raphael, 1987; Brener *et al.*, 2003). A natural question is therefore whether self-reported data is appropriate for studying cessation dynamics and intervention effects. Although details on the field of psychometrics is outside the scope of this dissertation (see Rust and Golmbok, 1999 and Urbina, 2011), analysis of self-reported data in this work is motivated by the following: self-reports have been shown to be informative and representative of actual psychological and physiological processes (Brener *et al.*, 2003; Patrick *et al.*, 1994). Haley *et al.* (1983), for example, demonstrated that self-reported smoking behavior strongly correlated with biochemical markers of smoking activity. Furthermore, the surveys in the McCarthy *et al.* (2008b) study drew from the Wisconsin Smoking Withdrawal Scale (WSWS; Welsch *et al.*, 1999) and the Positive and Negative Affect Schedule (PANAS; Watson *et al.*, 1988; Watson and Clark, 1999). These are well-established measures that have been rigorously evaluated for their relevance in assessing withdrawal and emotional constructs. Castro *et al.* (2011), for example, examines the predictive ability of the WSWS. Similarly, Watson and Clark (1999) details the construct validity (whether the survey appropriately measures the psychological factors it claims to measure) and external validity (generalizability of the data collected with this measure) of the PANAS (Shadish *et al.*, 2002; Watson and Clark, 1999). These considerations suggest models estimated from self-reported behaviors and related constructs can describe psychological phenomena.

Furthermore, more objective metrics of smoking activity are not necessarily the most appropriate variables to consider in this work. Notably, biochemical measurements of nicotine metabolites are poor indicators of smoking activity if the patient is being treated with nicotine-based therapies. As five of the seven first-line tobacco dependence pharmacotherapies involve low dose nicotine delivery (Tobacco Use and

Dependence Guideline Panel, 2008; Rennard *et al.*, 2014), serum nicotine level or similar indicators are impractical behavioral metrics. Similarly, carbon monoxide levels in expired air varies significantly over time and can be a function of environmental circumstances unrelated to cigarette smoking (Fortmann *et al.*, 1984), and is similarly an impractical measurement of cessation-induced behavioral dynamics.

Estimating dynamic models with self-reported CPD, $Craving$, and $Stress$ levels is further motivated by the intervention design component of this work. A smoker smokes based on an implicit or explicitly-acknowledged desire to do so. Therefore, estimating dynamic models estimated using self-reported data on this level is more in line with the "simplest useful model" aspect of designing a predictive controller (Rossiter, 2003). In other words, modeling smoking behavior change with a first-principles approach is suboptimal, given that these behaviors are the result of a confluence of biochemical, physiological, learned habitual factors, and more.

2.3.4 System Identification

System identification encompasses a set of principles and techniques to obtain estimated dynamical systems models that are useful in many end-use applications. A formal system identification procedure involves four components:

1. Design and implementation of an experiment to study a dynamic system of interest under various operating conditions. The result should be an informative data set conducive to estimation of high-fidelity dynamical systems models.

2. Data pre-processing to remove drifts and trends, address high-frequency disturbances above the frequency of interest, and deal with anomalies and missing data. The resulting data set then informs model structure selection.

3. Parameter estimation for a user-defined equation structure.

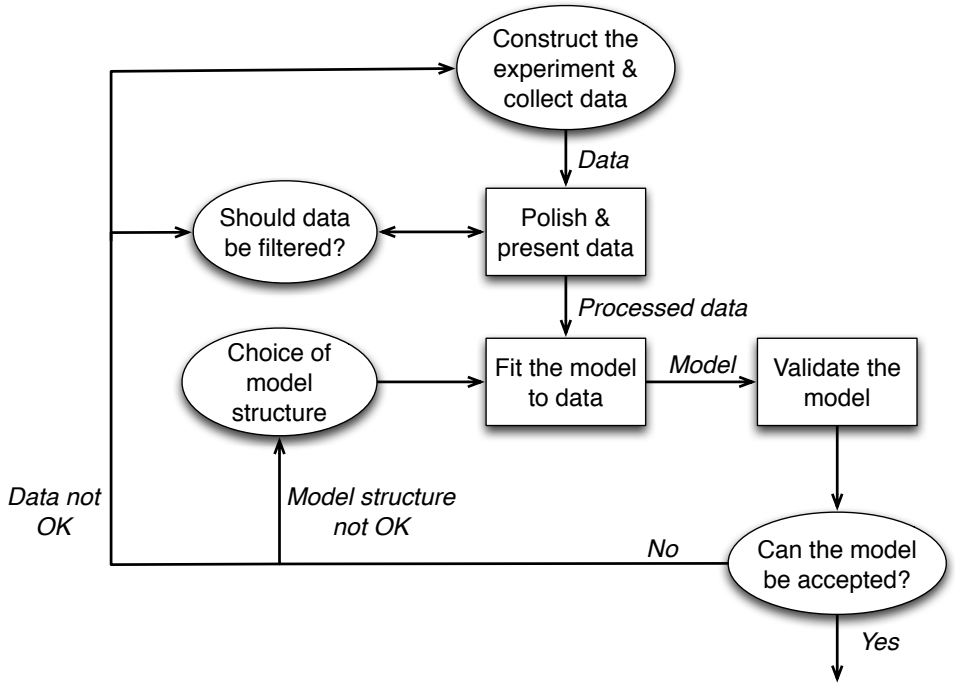

Figure 2.1: Flow Chart of the Iterative System Identification Process (Ljung, 1999).

4. Model validation. This can be broad-based (e.g., white residuals) or focused on a particular model end-use (e.g., control-relevant identification).

These steps should actually be components of an iterative procedure, as depicted in the flow chart in Fig. 2.1 (Ljung, 1999).

Detailed later in this chapter, Model estimation in this dissertation primarily involves secondary data analysis of ILD collected in the McCarthy *et al.* (2008b) clinical trial, as will be detailed later in this chapter. Consequently, this work relies most heavily on the parameter estimation aspects of system identification. More specifically, a prediction-error estimation approach is primarily employed to estimate continuous-time models from sampled data. The remainder of this subsection provides details on this approach.

Estimating Continuous-Time Models with Prediction-Error Methods

Continuous-time dynamic models offer a number of distinct advantages generally (Garnier *et al.*, 2008b; Garnier and Young, 2014) and in this behavioral science context specifically (Timms *et al.*, 2014a). This work primarily estimates continuous-time transfer functions with a two-step indirect approach: first, a discrete-time model is estimated from sampled data; next, the discrete-time model is transformed into continuous-time form. Indirect continuous-time identification benefits from the reliability and statistical efficiency of discrete-time model estimation methods (Garnier *et al.*, 2008b). Direct continuous-time model identification techniques are detailed in Garnier *et al.* (2008b) and Young *et al.* (2008), but lie outside the scope of this dissertation. The general estimation method is described in the following.

Consider an AutoRegressive with eXternal input (ARX) discrete-time model of the general form:

$$A(q^{-1}, \boldsymbol{\theta})y(t) = B(q^{-1}, \boldsymbol{\theta})u(t) + e(t) \tag{2.2}$$

where t represents a specific instance in time (a sample) and q^{-1} is the time-domain backshift operator, i.e., $q^{-1}u(t) = u(t-1)$. $A(q^{-1}, \boldsymbol{\theta})$ is the autoregressive polynomial of order n_a and $B(q^{-1}, \boldsymbol{\theta})$ is the external input polynomial of order n_b:

$$A(q^{-1}, \boldsymbol{\theta}) = 1 + a_1 q^{-1} + a_2 q^{-2} + \cdots + a_{n_a} q^{-n_a} \tag{2.3}$$

$$B(q^{-1}, \boldsymbol{\theta}) = b_1 q^{-1} + b_2 q^{-2} + \cdots + b_{n_b} q^{-n_b} \tag{2.4}$$

Equation 2.2 can be rewritten as,

$$y(t) = \boldsymbol{\varphi}^T(t)\boldsymbol{\theta} + e(t) \tag{2.5}$$

where the regressor and vector of parameters are, respectively,

$$\boldsymbol{\varphi}(t) = [-y(t-1) \ \ldots \ -y(t-n_a) \ u(t-1) \ \ldots \ u(t-n_b)]^T \tag{2.6}$$

$$\boldsymbol{\theta} = [a_1 \ \ldots \ a_{n_a} \ b_1 \ \ldots b_{n_b}]^T \tag{2.7}$$

A classical least squares regression problem would analytically estimate $\boldsymbol{\theta}$ according to the expression in equation 2.8:

$$\hat{\boldsymbol{\theta}}_{LS} = \left(\frac{1}{N} \sum_{t=1}^{N} \boldsymbol{\varphi}(t) \boldsymbol{\varphi}^T(t) \right)^{-1} \frac{1}{N} \sum_{t=1}^{N} \boldsymbol{\varphi}(t) y(t) \tag{2.8}$$

However, $\hat{\boldsymbol{\theta}}_{LS}$ may be inconsistently estimated when the elements in $\boldsymbol{\varphi}(t)$ are measurements that are contaminated by noise—when $e(t)$ in equation 2.5 is not white noise—as is the typical case (Garnier *et al.*, 2008b). To manage this, the least squares approach is modified according to instrumental variables methods (where the normal equations associated with the least squares estimate are altered) or prediction error methods (PEM; Ljung, 1999). The latter set of methods are employed in this work.

Generally, PEM seeks to model the noise in the measured samples as well. In other words, the model outputs are considered to be stochastic, as opposed to deterministic (Ljung, 1999). PEM encompasses a family of estimation approaches that generally seek to estimate $\hat{\boldsymbol{\theta}}$ by minimizing a weighted norm of the prediction error:

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}} V_N(\boldsymbol{\theta}) \tag{2.9}$$

where V_N is a cost function that is a scalar value (Garnier *et al.*, 2008b; Ljung, 1999; The MathWorks, 2014e):

$$V_N(\boldsymbol{\theta}) = \frac{1}{N} \sum_{t=1}^{N} \ell(L(q^{-1})\varepsilon(t, \boldsymbol{\theta})) \tag{2.10}$$

ε is the vector of prediction errors. Specifically, ε is the difference between the observed data, $y(t)$, and the model predictions, $\hat{y}(t|t-1, \boldsymbol{\theta})$ at each sample, t:

$$\varepsilon(t, \boldsymbol{\theta}) = y(t) - \hat{y}(t|t-1, \boldsymbol{\theta}), \ \ t = 1, \ldots, N \tag{2.11}$$

$\ell(\cdot)$ in equation 2.10 denotes a scalar-valued function that measures the size or norm of the expression it encloses (Garnier *et al.*, 2008b; Ljung, 1999); $L(q^{-1})$ is a linear filter that can weigh certain properties of the model (e.g., filter out high-frequency disturbance effects; Ljung, 1999).

With the task in equation 2.9 in mind, consider the general model structure:

$$y(t) = G(q^{-1}, \boldsymbol{\theta})u(t) + H(q^{-1}, \boldsymbol{\theta})e(t) \tag{2.12}$$

where $G(q^{-1}, \boldsymbol{\theta})$ is the u to y transfer function and $H(q^{-1}, \boldsymbol{\theta})$ is the e to y transfer function. Equating $e(t)$ and $\varepsilon(t)$ and substituting $e(t)$ into equation 2.23 gives:

$$y(t) = G(q^{-1}, \boldsymbol{\theta})u(t) + H(q^{-1}, \boldsymbol{\theta})(y(t) - \hat{y}(t|t-1, \boldsymbol{\theta})) \tag{2.13}$$

Solving for $\hat{y}(t|t-1, \boldsymbol{\theta})$ and rearranging gives the optimal mean square predictor:

$$\hat{y}(t|t-1, \boldsymbol{\theta}) = H^{-1}(q^{-1}, \boldsymbol{\theta})G(q^{-1}, \boldsymbol{\theta})u(t) + (I - H^{-1}(q^{-1}, \boldsymbol{\theta}))y(t) \tag{2.14}$$

If equation 2.14 is a linear function with respect to the unknown parameters, and equation 2.10 consists of a relatively simple cost function, a closed form solution can be obtained. For example, when equation 2.14 can be written as,

$$\hat{y}(t|t-1, \boldsymbol{\theta}) = \boldsymbol{\varphi}^T(t)\boldsymbol{\theta} \tag{2.15}$$

and V_N is quadratic,

$$V_N(\boldsymbol{\theta}) = \frac{1}{N}\sum_{t=1}^{N}\varepsilon^2(t, \boldsymbol{\theta}) = \frac{1}{N}\sum_{t=1}^{N}(y(t) - \boldsymbol{\varphi}^T(t)\boldsymbol{\theta})^2 \tag{2.16}$$

estimation of $\hat{\boldsymbol{\theta}}$ can follow computations from classical regression procedures, notably, equation 2.8 (Garnier *et al.*, 2008b; Ljung, 1999). For more complex cost functions or when nonlinearities are present in equation 2.14, estimation of $\hat{\boldsymbol{\theta}}$ relies on a numerical search algorithm (Ljung, 1999). Common numerical search algorithms include subspace Gauss-Newton approach, Levenberg-Marquart method, and steepest-descent gradient method (Ljung, 1999; The MathWorks, 2014f).

2.4 Understanding Statistical Mediation & Self-Regulation Within Smoking Behavior Change

The remainder of this section is a copy of the Timms *et al.* (2014a) publication. This manuscript details development of behavior change models of statistical mediation and self-regulation in a dynamical systems sense. (Note, Timms *et al.* (2014a) refers to the CPD signal as $Cigsmked$.) Timms *et al.* (2014c) and Timms *et al.* (2014b) offer a more tutorial explanation of the mediation and self-regulation modeling methodology. Extensions of the modeling ideas described in Timms *et al.* (2014a), Timms *et al.* (2014b), and Timms *et al.* (2014c) are documented in Section 2.5.

2.4.1 Continuous-Time System Identification of a Smoking Cessation Intervention

Cigarette smoking is a major global public health issue. Approximately 10 million annual global deaths are expected to result from smoking by 2020 (Fish and Bartholomew, 2007), and the global smoking population is expected to surpass 1.7 billion by 2025 (Erhardt, 2009). In the U.S., cigarette smoking is the leading cause of premature death and $157B in economic loss is attributed to tobacco use annually (Centers for Disease Control and Prevention, 2010; Killeen, 2011). 40 years of decreases in smoking rates have recently stalled in the U.S., where one in five adults is an active smoker (Centers for Disease Control and Prevention, 2012). This smoking rate persists despite the fact that nearly 70% of smokers have expressed a desire to quit (Center for Disease Control and Prevention, 2011). Largely due to the chronic, relapsing nature of cigarette smoking (Tobacco Use and Dependence Guideline Panel, 2008), over 88% of attempts to quit smoking fail (American Cancer Society, 2012).

Interventions play an important role in smoking prevention and cessation. Generally speaking, interventions for behavioral health disorders seek to reduce unhealthy

behaviors and promote healthy ones through prevention or treatment, and are used to address many public health concerns in addition to smoking such as other substance abuse, obesity, sexually transmitted infections, and cancer screening, and can be pharmacological or behavioral in nature (Collins, 2012; Baker *et al.*, 2011). Traditionally, these interventions are "fixed," meaning they are not systematically operationalized, and the composition and dosage of an intervention component is given to all individuals receiving the intervention and do not vary over time (Collins *et al.*, 2004). The effectiveness of existing fixed smoking cessation interventions is limited. Counseling alone, for example, has a reported success rate below 15% (Tobacco Use and Dependence Guideline Panel, 2008; Fish and Bartholomew, 2007). Pharmacological interventions (e.g., nicotine replacement therapies such as Nicorette®) have individual one-year abstinence rates below 35% (Fish and Bartholomew, 2007), which may be lower at longer term follow-up (Irvin and Brandon, 2000; Irvin *et al.*, 2003). Such low success rates are particularly troubling given the gravity of cigarette smoking as a public health issue.

To address the limitations of fixed interventions, recent efforts in behavioral health have centered around development of so-called "adaptive" interventions, where treatment components and dosage vary according to participant response (Collins *et al.*, 2004). These interventions consist of closed-loop dynamical systems and may be more effective behavioral health interventions as they essentially seek to optimally adapt to the changing needs of a patient (Rivera *et al.*, 2007; Nandola and Rivera, 2013). Control systems engineering principles offer an appealing framework for developing algorithms that implement these optimized, time-varying smoking cessation interventions. However, the impact of using control engineering concepts in the design of time-varying smoking cessation interventions is tied to the reliability of the smoking behavior change models upon which the algorithms are based (Nandola and Rivera,

2013; Rivera, 2012; Riley *et al.*, 2011).

The development of reliable models has been greatly enhanced by increased access to intensive longitudinal data (ILD). ILD in behavioral settings is loosely defined as quantitative or qualitative measurements recorded frequently over time, and is more readily available due to increasing use of mobile and computerized technologies in behavioral trials (Walls and Schafer, 2006). ILD facilitates the dynamic modeling of behavior, and generally provides a means for improved analysis of inter- and intra-individual variability (Collins, 2006). In the context of cigarette smoking, ILD offers an opportunity to study the dynamics of smoking behavior change (e.g., daily smoking rate, average craving level) during a quit attempt. Whereas traditional quantitative modeling methodologies from behavioral science (e.g., structural equation models, SEMs) are static in nature (Bollen, 1989), dynamical systems modeling and system identification offer a framework for more comprehensive characterization of dynamic behavioral relationships and how smoking cessation interventions affect these dynamics (Timms *et al.*, 2012). Recently, similar models have been used for improved evaluation of gestational weight gain and fibromyalgia interventions; these models also offer an appealing basis for development of optimized, time-varying interventions that draw from control systems engineering principles (Dong *et al.*, 2013; Nandola and Rivera, 2013; Rivera, 2012; Deshpande *et al.*, 2011). In terms of smoking cessation, an intervention based in controller design methods could feature feedback in order to systematically assign medication dosages based on patient reports of withdrawal symptoms, for example (Timms *et al.*, 2013). Given the gravity of cigarette smoking as a public health issue and the modest effectiveness of even the most efficacious treatments available, an improved ability to inform and evaluate behavioral health interventions warrants development of dynamic models of smoking cessation.

In this article, smoking cessation is described as a process of behavior change, and

this process is represented via continuous-time models. A continuous-time approach is particularly appealing in behavioral health settings (Timms *et al.*, 2012; Rivera, 2012). Notably, for low order dynamics (which appears to be the case for many dynamic behaviors; Timms *et al.* 2012, 2014c; Deshpande *et al.* 2011; Rivera 2012), continuous-time models estimated from discrete-time data result in parsimonious expressions through which important dynamic features are more easily discerned. Consequently, insight into behavioral phenomena and intervention effects may be more easily interpretable with continuous-time models. As will be demonstrated, continuous-time models that capture inverse response in craving during a quit attempt can be easily identified with a right half plane zero term, the estimated value of which may shed light on underlying dynamic phenomena (Timms *et al.*, 2012, 2013). Similarly, the parameters in continuous-time models are more meaningful in terms of understanding a process, as such models more transparently describe the continuous nature of actual physical systems of interest. An *a priori* understanding of a system can also be more easily preserved with continuous-time models, whereas discrete-time models of a second order system, for example, may introduce additional parameters due to sampling. Furthermore, discrete-time models estimated at fixed sampling rates may not be representative of process dynamics observed under different sampling rates (Garnier *et al.*, 2008b). This may be an important consideration given the range of time scales of interest to behavioral scientists (e.g., short term, long term, non-standard time periods such as pubertal time scales; Collins *et al.* 2004). Additionally, missing data and irregular sampling intervals are characteristic of self-reported behavioral health experiments, whether intentional or not; this supports the appeal of a continuous-time modeling approach, which will not have parameters that are a function of the sampling time, and therefore inherently manage the issue of non-ideal data measurement. Consequently, continuous-time models of discrete data collected

at a sufficiently fast sampling rate can produce models that accurately represent the overall dynamics of behavioral phenomena, while avoiding challenges associated with missing data and inconsistent sampling intervals in discrete-time modeling (Garnier *et al.*, 2008b).

Behavioral scientists often rely on behavior change theories to guide intervention design, evaluation, and delivery (Riley *et al.*, 2011). In the context of cigarette smoking, the concepts of statistical mediation and self-regulation have been of particular interest (Timms *et al.*, 2012, 2014c; McCarthy *et al.*, 2008a; Piper *et al.*, 2008; Velicer *et al.*, 1992; Walls and Rivera, 2009). Statistical mediation is a modeling paradigm central to the social and behavioral sciences, describing a multivariate causal relationship in which an independent variable affects a mediator variable and an outcome variable, with the mediator also affecting the outcome (MacKinnon, 2008). Self-regulation theory within smoking considers a process in which nicotine levels, behavioral state, or emotional state set points are regulated by smoking activity (Velicer *et al.*, 1992; Walls and Rivera, 2009; Solomon and Corbit, 1974; Timms *et al.*, 2013, 2014c). However, the utility of models that describe these behavior change theories for the purposes of development of optimized, time-varying interventions has been limited; this is largely a consequence of the static nature of traditional behavioral science models and the difficulties historically associated with intensive collection of behavioral data (Riley *et al.*, 2011). This article employs a modeling framework that leverages ILD and continuous-time system identification in order to describe smoking cessation as a mediational and self-regulatory process.

This section is organized as follows. First, a clinical trial of bupropion and counseling as aids to smoking cessation is outlined; this study was conducted at the University of Wisconsin Center for Tobacco Research and Intervention (UW-CTRI) and funded by the Transdisciplinary Tobacco Use Research Centers (TTURC; McCarthy

et al. 2008b). Behavioral signals for treatment group averages and two single subject examples from the clinical trial are then presented. The features of the signals of interest in this article—*Craving*, average daily craving level, and *Cigsmked*, total daily smoking—are discussed in general terms. Statistical mediation is first presented in conceptual terms and the connection to dynamic model development is then outlined. The iterative procedure used for model estimation is then described and the resulting empirical dynamic mediation models are discussed. Next, self-regulatory smoking cessation models are presented, estimated, and compared. In examination of both the mediation and self-regulation models, parameter estimates and model simulations are analyzed for treatment group averages. Following this, examples of single subject models are briefly discussed. Finally, conclusions and recommendations are presented.

Smoking Cessation Intervention Overview

Dynamic models are obtained in a secondary analysis of a TTURC-funded study conducted by the UW-CTRI. In this double-blinded, placebo-controlled, randomized clinical trial, 101 subjects received both active bupropion and counseling as treatment (the "AC" group), 101 received active bupropion and no counseling ("ANc"), 100 received a placebo and counseling ("PC"), and 101 received a placebo and no counseling ("PNc"). Participants receiving bupropion took 150 mg per day starting one week prior to the quit date and 300 mg per day from four days prior to quit to eight weeks post-quit. Bupropion SR (Zyban SR) is commonly prescribed as a smoking cessation treatment (Fish and Bartholomew, 2007; Tobacco Use and Dependence Guideline Panel, 2008), although the exact mechanism that makes it an effective smoking treatment is debated (Horst and Preskorn, 1998; Warner and Shoaib, 2005). Generally, bupropion is thought to interfere with nicotine dependence mechanisms

(Warner and Shoaib, 2005), and has been shown to alleviate withdrawal symptoms, including craving (McCarthy *et al.*, 2008b). In lieu of active bupropion, the PC and PNc groups took placebo medication. Subjects receiving counseling completed two pre-quit counseling sessions, one quit-date session, and five sessions over the following four weeks post-quit. Sessions focused on preparation, coping, motivation, and relapse prevention. In lieu of counseling, the ANc and PNc groups spoke with staff about medication use adherence and received general encouragement (McCarthy *et al.*, 2008b).

For the two weeks prior to and four weeks immediately following the target quit date, participants were instructed to complete assessments through personal digital assistants (PDAs) each day immediately after waking up, before going to bed, and four to seven times throughout the day as prompted by the PDA at pseudo-random intervals. These self-reports generally collected data on smoking frequency and withdrawal symptoms. Although data from each set of reports or a combination of reports could be used to model the smoking process, the Evening Report (ER) is the focus of this article's efforts. The ER featured questions on a 10 point Likert scale covering topics such as withdrawal symptoms (McCarthy *et al.*, 2008b), positive affect (generally, the degree to which an individual feels enthusiastic and alert), negative affect (generally, the degree to which an individual feels anger, disgust, guilt, fear, and nervousness; Watson *et al.* 1988), and motivation to abstain. Table 2.1 provides a selection of items from the ER (McCarthy *et al.*, 2008b). The relationship between *Craving* and *Cigsmked* variables is the focus of this article, as was done in a statistical study of the same ILD by McCarthy *et al.* (2008a). *Craving* is defined as the sum of *Urge*, *Cigonmind*, *Thinksmk*, and *Bother*. *Cigsmked* is the total number of cigarettes smoked per day.

Both nomothetic (group level) and idiographic (single participant) models are of

Table 2.1: Representative Subset of Questions from the Evening Report.

Code	Question	Scale
Urge	Since last ER on average—Bothered by urges?	1-11, No!!...Yes!!
Cigonmind	Since last ER on average—Cigarettes on my mind?	1-11, No!!...Yes!!
Thinksmk	Since last ER on average—Thinking about smoking a lot?	1-11, No!!...Yes!!
Bother	Since last ER on average—Bothered by desire to smoke?	1-11, No!!...Yes!!
Enthus	Since last ER on average—Enthusiastic?	1-11, No!!...Yes!!
Tense	Since last ER on average—Tense or anxious?	1-11, No!!...Yes!!
Sad	Since last ER on average—Sad or depressed?	1-11, No!!...Yes!!
Anger	Since last ER on average—Bothered by anger/irritability?	1-11, No!!...Yes!!
Excellent	Since last ER on average—Concentration was excellent?	1-11, No!!...Yes!!
Happy	Since last ER on average—Happy and content?	1-11, No!!...Yes!!
Food	Since last ER on average—Thinking about food a lot?	1-11, No!!...Yes!!
Confidence	Since last ER on average—Confidence in ability to quit?	1-11, Low!!...High!!
Motive	Since last ER on average—Motivation to quit/stay quit?	1-11, Low!!...High!!
Cigsmked	Total number of cigarettes smoked since the last ER?	0-99

general interest to behavioral scientists. In this section, treatment group average models are the primary focus. To produce the group average signal, each report item was averaged across all members in a group for each relative day over the week prior to and four weeks immediately following the target quit date. This filtering that occurs by averaging the data across all single subject data points in a group, and putting the time series in deviation variable form, was the only data pre-processing done prior to group average model estimation. While the continuous-time approach employed here can effectively model data with missing samples or non-constant sampling, missing ER data for the two single subject examples was interpolated for straightforward use of standard MATLAB estimation routines. Interpolation consisted of averaging adjacent measured values or extending the adjacent measured value to the appropriate boundary. For the single subject example from the AC group, eight days of data

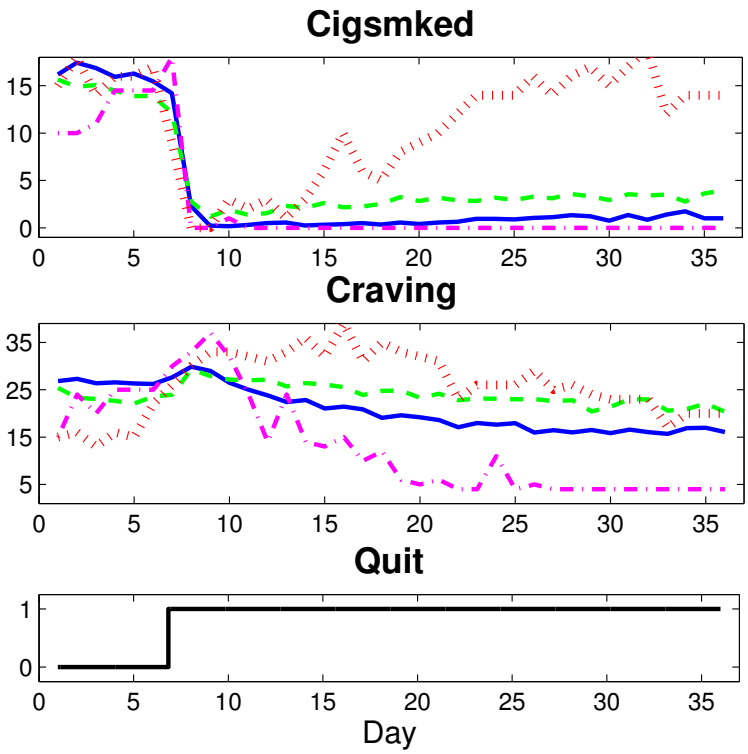

Figure 2.2: Plots of Two group Average (Solid Blue, AC; Dashed Green, PNc) and Two Single Subject (Dash-Dot Magenta, AC; Dotted Red, PNc) Data Sets.

points were imputed. Seven days of data points were imputed for the single subject example from the PNc group. Although single subject data sets are often noisier than corresponding group average data, no filtering was done on these data prior to idiographic model estimation. Fig. 2.2 depicts the *Craving* and *Cigsmked* ILD for two group averages (solid blue, AC; dashed green, PNc) and two single subject examples (dash-dot magenta, AC; dotted red, PNc).

As seen in Fig. 2.2, the group average *Craving* signals feature quit-induced inverse response. With a continuous-time modeling approach, *a priori* knowledge that the groups' *Craving* signals features a right half plane zero is more easily preserved (Garnier *et al.*, 2008b). The group average *Cigsmked* signals feature a dramatic quit-day drop, followed by a relatively small and slow resumption of smoking. The single subject data sets display greater variability. In Fig. 2.2, the PNc single subject

does not feature a net reduction in craving. The AC subject has little resumption in smoking—reflecting a successful quit attempt—while the PNc subject features significant resumption to approximately pre-quit levels.

Statistical Mediation Modeling

The concept of statistical mediation is a prominent model of change in the social and behavioral sciences. As previously described, mediation defines a causal relationship in which an independent variable, X, affects a mediator, M, and an outcome, Y, with M also contributing to Y (MacKinnon, 2008). Behavioral scientists use path diagrams to depict this type of process (MacKinnon, 2008; Bollen, 1989). A mediational model path diagram—not to be confused with a block diagram—is depicted in Fig. 2.3(a): a, b, and c' represent gains for the X to M, M to Y, and X to Y pathways, respectively (MacKinnon, 2008). Structural Equation Model (SEM) representations of mediation are found in Equations 2.17 and 2.18.

$$M = \beta_{0_1} + a\,X + e_1 \tag{2.17}$$

$$Y = \beta_{0_2} + c'\,X + b\,M + e_2 \tag{2.18}$$

Historically, Equations 2.17 and 2.18 have characterized mediation thought to be captured within cross-sectional studies; X typically codes the presence or absence of an intervention and M and Y data are collected at a small number of time points (MacKinnon, 2008). A dynamical framework is developed in this manuscript according to a more general definition of mediation described in Collins *et al.* (1998). Collins *et al.* (1998) underscores a temporal relationship between X, M, and Y, where a change in an independent variable at some time is said to result in lagged changes in the mediator and outcome (Collins *et al.*, 1998). The SEMs in Equations 2.17 and 2.18 still apply under the Collins *et al.* (1998) definition, but the variables are a function

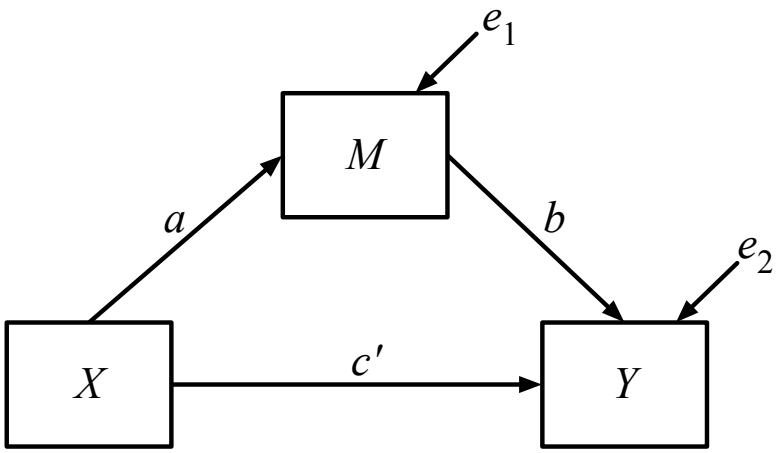

(a) Path Diagram, Used in Structural Equation Modeling (SEM) to Describe the Relationship Between Variables, for a Classic Mediational Model.

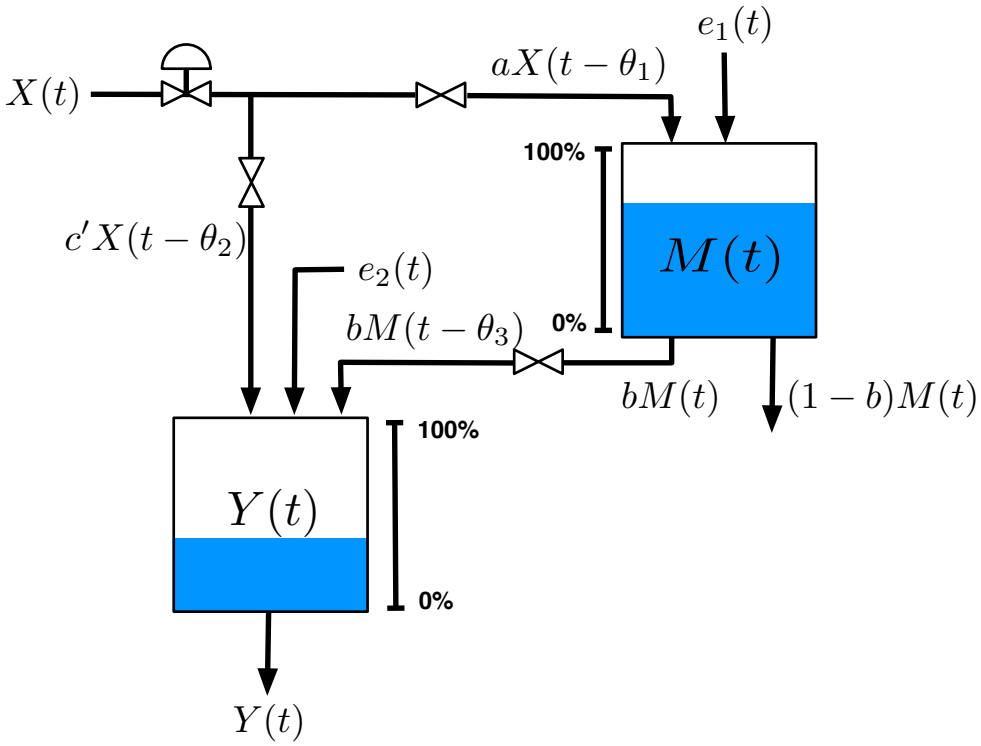

(b) Fluid Analogy for Mediated Behavior Change Developed from Production Inventory Management Models in Supply Chains.

Figure 2.3: Path Diagram Describing Statistical Mediation and an Initial Fluid Analogy.

of specific, discrete times.

Prior work by Navarro-Barrientos *et al.* (2011) established how path diagrams in SEM correspond to steady-state process models; from these, fluid analogies can be constructed which lead to dynamical systems amenable to estimation via system identification methods. Drawing from techniques used in production inventory management in supply chains, fluid analogies describe dynamic behaviors according to a structural relationship defined by a behavioral model (Schwartz *et al.*, 2006). The fluid analogy in Fig. 2.3(b) depicts a physical system analogous to behavior change according to a mediational mechanism. Here, the independent variable corresponds to the exogenous input to the system, and the endogenous quantities in the path diagram (M and Y in Fig. 2.3(a)) are represented as inventories. Dynamic, mediated behavior change models are obtained when each inventory is considered in the context of a general conservation principle:

$$\text{Accumulation} \;=\; \text{Sum of Inflows} \;-\; \text{Sum of Outflows} \tag{2.19}$$

A system of basic, first order differential equations results from the application of Equation 2.19 to Fig. 2.3(b):

$$\tau_1 \frac{dM}{dt} \;=\; aX(t-\theta_1) - M(t) + e_1(t) \tag{2.20}$$

$$\tau_2 \frac{dY}{dt} \;=\; c'X(t-\theta_2) + bM(t-\theta_3) - Y(t) + e_2(t) \tag{2.21}$$

where the derivative terms describe the changing levels of the inventory, a, b, and c' are the system gains, τ_1 and τ_2 are the inventory time constants, and θ_1, θ_2, and θ_3 are the time delays. It can be shown that at steady-state, Equations 2.20 and 2.21 simplify to the structural models in Equations 2.17 and 2.18.

Higher order differential equations could be used to describe more complex dynamic behavior while still adhering to relationships depicted in Fig. 2.3(b). While

this simple fluid analogy reflects structural relationships defined by Equations 2.17 and 2.18, the resulting differential equations are relatively restrictive. Specifically, the outcome inventory dynamics are bound by a single time constant, despite the fact that the inventory accepts contributions from both the independent variable input and the mediator inventory outlet. This restriction is not necessary in a behavioral setting. Fig. 2.4(a) is a less restrictive fluid analogy that describes mediated behavior change. In this analogy, each pathway from Fig. 2.3(a) is represented by an inventory, each with its own characteristic dynamics. Fig. 2.4(b) is the corresponding block diagram and highlights the fact that Y is the result of two processes, where Y_D accounts for the outcome change that is a direct result of the input variable change, and Y_I accounts for the outcome change that is an indirect result of the input variable change via the mediator variable change. Fig. 2.4(b) also highlights the parallel-cascade nature of time-varying behaviors in a mediational relationship. Equations 2.22 and 2.23 are the corresponding Laplace-domain models:

$$M(s) = P_a(s)\, X(s) + d_1(s) \tag{2.22}$$

$$Y(s) = P_{c'}(s)\, X(s) + P_b(s)\, M(s) + d_2(s) \tag{2.23}$$

In accordance with McCarthy *et al.* (2008a), this section primarily treats *Craving* and *Cigsmked* as the mediator and outcome, respectively. The independent variable input, *Quit*, is modeled as a unit step occurring on the quit date and corresponds to a transition from not attempting to quit smoking to attempting to quit. d_1 and d_2 in Equations 2.22 and 2.23 represent process disturbances, as opposed to measurement noise. In this framework, they represent un-modeled factors that influence the mediator and outcome. In this context, d_1 represents factors other than the initiation of a quit attempt that contribute to, or mitigate, *Craving*, such as negative or positive life events (i.e., changes in *Stress*); d_2 represents factors other than the initiation of a

quit attempt or changing *Craving* levels that influence *Cigsmked*, such as the accessibility or inaccessibility of cigarettes. If these process disturbances are measured and uncorrelated with the inputs to $P_a(s)$, $P_b(s)$, and $P_{c'}(s)$, they could be explicitly modeled (Ljung, 1999). In behavioral health settings, explicitly modeling disturbances is challenging, as additional influences that are truly independent and exogenous are often not obvious. Consequently, accurate measurement and estimation of reliable models presents significant practical and related estimation issues. In the following, it is assumed that process disturbances are uncorrelated with the inputs and not of significant influence. The group averaging likely effectively filters out un-controlled disturbances, suggesting these assumptions are reasonable. Ultimately, reliable identification and characterization of possible process disturbances would require novel identification experiments. With novel clinical trial data, measurements of hypothesized disturbances could be included in estimation of Equations 2.22 and 2.23 and validated (e.g., through cross-validation); future validation of a time-varying disturbance could significantly contribute to development of an improved—engineering-based or otherwise—smoking intervention, as it would ultimately offer tobacco treatment practitioners an additional degree-of-freedom on which to intervene.

In fitting the Laplace-domain models in Equations 2.22 and 2.23 to the *Craving* and *Cigsmked* signals, a prediction-error approach is used to estimate continuous-time linear models from sampled data (Ljung, 2009). Model estimation initially employed the lowest order equation structure possible (a gain), as preliminary visual inspection of the data indicated low order dynamic phenomena; transfer function structures with gradually increased complexity were evaluated as necessary according to the following iterative estimation and validation procedure:

(1) Estimation of $P_a(s)$ as a single-input single-output (SISO) system with *Quit* as the input and *Craving* as the output according to a given low order transfer

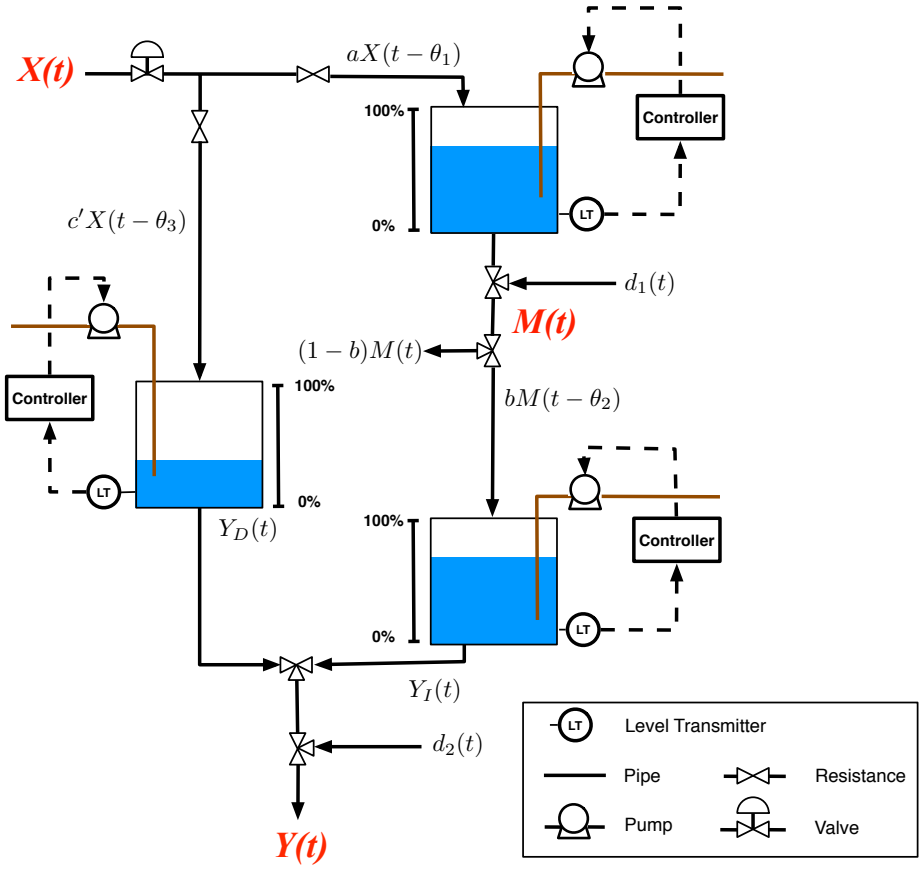

(a) Generalized Fluid Analogy for a Mediated Behavioral Intervention Developed from Production Inventory Management Models in Supply Chains.

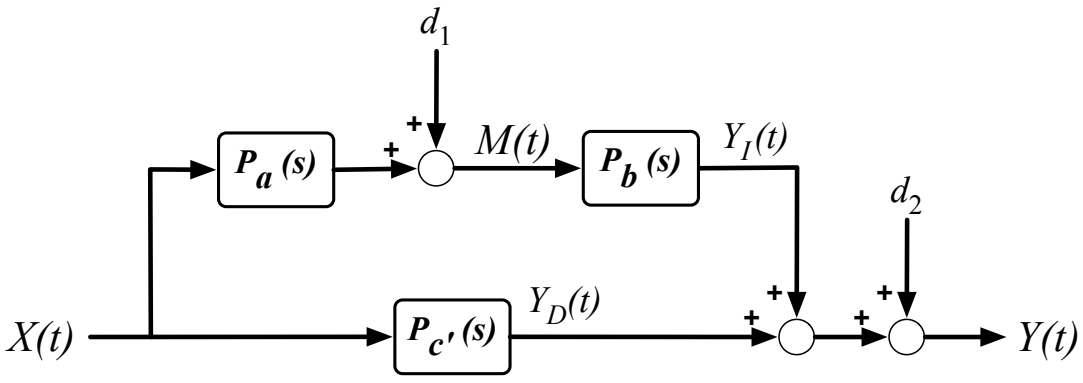

(b) Block diagram of statistical mediation.

Figure 2.4: Generalized Fluid Analogy and Block Diagram Describing Dynamic Mediation.

function structure.

(2) Simultaneous estimation of $P_b(s)$ and $P_{c'}(s)$ as a multi-input single-output (MISO) system with *Craving* and *Quit* as the inputs and *Cigsmked* as the output according to given low order transfer function structures for $P_b(s)$ and $P_{c'}(s)$.

(3) Simulation of the *Craving* and *Cigsmked* responses to *Quit* according to the estimated $P_a(s)$, $P_b(s)$ and $P_{c'}(s)$ expressions.

(4) Evaluation of *Craving* and *Cigsmked* goodness-of-fit on a 0 to 100% scale, calculated according to the following criterion:

$$Fit\ [\%] = 100\left(1 - \frac{||y(t) - \tilde{y}(t)||_2}{||y(t) - \bar{y}||_2}\right) \tag{2.24}$$

where $y(t)$ is the data to which the model is fit, $\tilde{y}(t)$ is the simulated output, and $\bar{y}$ is the average of all y values.

This procedure was implemented in MATLAB through a custom graphical user interface (GUI) built for flexible model estimation (see Fig. A.1). Using the GUI, the four steps were repeated for different combinations of $P_a(s)$, $P_b(s)$, and $P_{c'}(s)$ transfer function structures with parameter estimation relying on the `pem` command from the System Identification Toolbox in MATLAB. To use this routine, the input and output data was defined as an `iddata` object and the structure of the model to be estimated was defined as an `idproc` model object. For specification of the `idproc` model structure, the *process models* notation was used in which a single-output continuous-time model transfer function structure is specified and can feature one to three poles, an integrator, a zero, and a time-delay (Ljung and Singh, 2012). The `idproc` and process models functionality employs an indirect continuous-time estimation approach

in which discrete-time estimation methods are first used before the resulting discrete-time representation is transformed into the equivalent continuous-time model. This two-step approach has the advantage of drawing from established discrete-time estimation methods to produce consistent and statistically efficient parameter estimates, and still results in continuous-time models that parsimoniously represent complex behaviors and can more easily be interpreted (Garnier *et al.*, 2008b). In the iterative four step procedure used here, a set of model estimates were ultimately identified as appropriate representations of the behavioral dynamics for each group average and the single subject examples according to goodness-of-fit values, a concern for model parsimony, evaluation of parameter realizability, and through simulation.

Table 2.2 contains the parameter estimates, settling times (in days), and goodness-of-fit values for the mediation models; ILD and model outputs are shown in Figures 2.5 and 2.6. The iterative estimation procedure's employment of low order structures where complexity was increased only as necessary, the associated high goodness-of-fit values, and the corresponding simulations, which appear to accurately model the dynamic features observed in the ILD, suggest the following transfer function structures adequately represent cessation as a $Craving$-mediated process:

$$P_a(s) = \frac{a\,(\tau_a s + 1)}{(\tau_1 s + 1)} \tag{2.25}$$

$$P_b(s) = \frac{b}{(\tau_3 s + 1)} \tag{2.26}$$

$$P_{c'}(s) = c' \tag{2.27}$$

The estimated models feature high fit percentages according to Equation 2.24 for the group averages. The low order of the structures indicate that over-parameterization is not taking place. In general, the high signal-to-noise ratios of the group average data sets are conducive to model estimation with high goodness-of-fit values, regardless of the transfer function structure. The lower mediator fit value for the PNc group

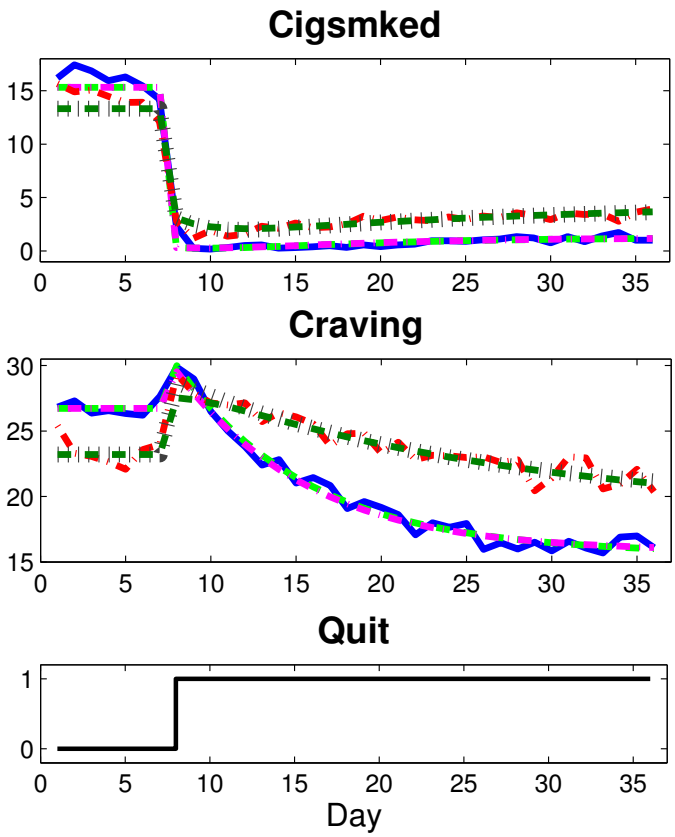

Figure 2.5: *Craving* and *Cigsmked* Data and Models for AC and PNc Group Averages (Solid Blue, AC Data; Dashed Light Green, AC Mediation Model; Dash-Dot Magenta, AC Self-Regulation Model; Dash-Dot Red, PNc Data; Dotted Brown, PNc Mediation Model; Dashed Dark Green, PNc Self-Regulation Model).

supports this assertion, as this group's *Craving* appears to have a lower signal-to-noise ratio than its counterparts.

The net decrease in *Craving* is greatest in the AC group, smallest in the PNc group, and follows a logical relationship to treatment condition (a is -11.10 for AC, -8.38 for ANc, -7.12 for PC, and -3.90 for PNc). As indicated by the negative system zeroes, the group average craving signals feature pronounced inverse response. It is known that a zero term in a dynamical systems model can result from two lower-order subprocesses in parallel. Because inverse response in *Craving* results from a unit step (*Quit*), it can be deduced that $P_a(s)$ reflects parallel subprocesses in competition with

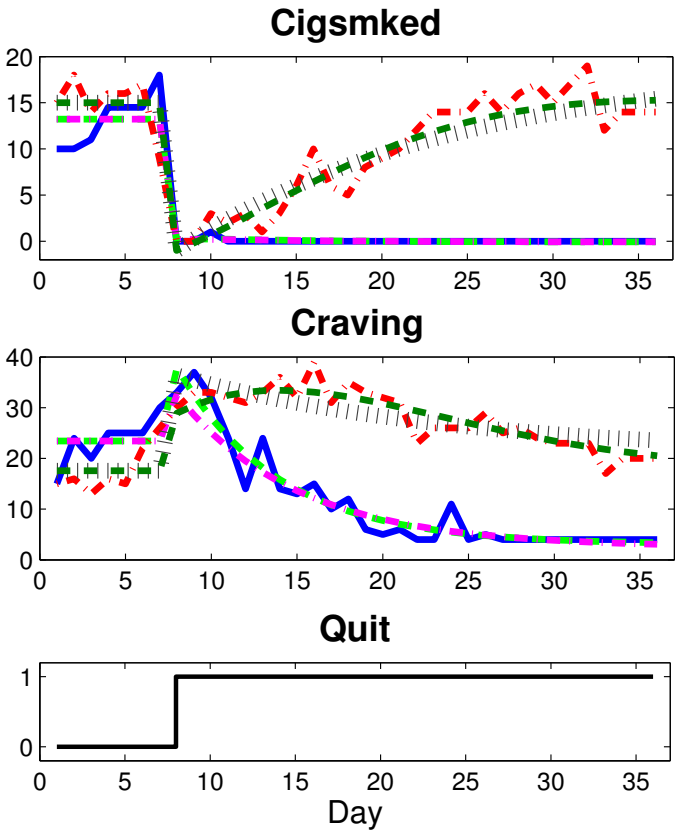

Figure 2.6: *Craving* and *Cigsmked* Data and Models for AC and PNc Single Subject Examples (Solid Blue, AC Data; Dashed Light Green, AC Mediation Model; Dash-Dot Magenta, AC Self-Regulation Model; Dash-Dot Red, PNc Data; Dotted Brown, PNc Mediation Model; Dashed Dark Green, PNc Self-Regulation Model).

each other. The first subprocess has a positive gain and faster speed of response than the second, negative-gained subprocess. The positive-gained subprocess corresponds to the immediate, quit-induced increase in *Craving* while the negative-gained subprocess corresponds to the post-quit settling of *Craving* to below baseline levels. For the case that $P_a(s)$ is described by the differential equation structure in Equation 2.25, $P_a(s) = P_{a_1}(s) + P_{a_2}(s)$; $P_{a_1}(s) = K_{a_1}$ and $P_{a_2}(s) = K_{a_2}/(\tau_1 s + 1)$, where the time constant is equal to that for the overall $P_a(s)$ function. It follows that $K_{a_1} = -a\tau_a/\tau_1$ and $K_{a_2} = a - K_{a_1}$, where a is the $P_a(s)$ gain, τ_a is the $P_a(s)$ zero, and τ_1 is the $P_a(s)$ time constant. This notion of competing parallel processes within the overall

Table 2.2: Mediation Model Parameter Estimates and Goodness-of-Fit Values.

Treatment	AC	ANc	PC	PNc	AC	PNc
Data Set	Avg	Avg	Avg	Avg	Sgl	Sgl
Mediator Fit [%]	87.77	78.88	77.80	64.72	69.44	44.80
Outcome Fit [%]	89.17	83.06	91.49	84.38	77.09	58.98
ab	0.94	1.96	1.45	1.17	-0.24	19.34
$ab + c'$	-14.07	-11.17	-12.05	-9.07	-13.29	3.55
Mediator Settling Time [Days]	35.69	35.91	35.82	35.90	26.34	33.87
Outcome Settling Time [Days]	34.56	35.26	35.60	35.29	10.64	33.86
a	-11.10	-8.38	-7.12	-3.90	-20.38	3.10
τ_a	-2.28	-4.60	-14.18	-24.21	-4.23	100.00
τ_1	7.74	10.99	18.34	17.13	6.01	16.47
b	-0.08	-0.23	-0.20	-0.30	0.01	6.25
τ_3	4.59	2.89	0.42	1.89	1.29	95.53
c'	-15.10	-13.13	-13.50	-10.42	-13.05	-15.99

smoking cessation process agrees with the observation that quitting smoking involves delayed and immediate gratification motives (executive and impulsive neurological processes, respectively) that compete during a quit attempt (Bickel *et al.*, 2007). Such insight into the nature of these underlying subprocesses highlights the utility of a continuous-time system identification approach, as the implications of the $P_a(s)$ transfer function's first order with zero structure were easily identified.

For the group average *Cigsmked* models, there is a dramatic quit-date drop in smoking followed by a relatively small and slow resumption. This dramatic quit-date smoking reduction is modeled by $P_{c'}(s)$. Considering the treatment group averages, the magnitude of the initial drop is largest for the AC group, a 15.01 cigarette per day decrease, and smallest for the PNc group, a 10.42 cigarette per day decrease. For each model corresponding to the parameter estimates in Table 2.2, the direct

contribution of the quit attempt to *Cigsmked* is immediate, and the Y_D signal in Fig. 2.4 acts as a step of magnitude c'. The resumption of smoking is modeled by the mediated pathway, specifically $P_b(s)$. For the group averages, the speed of resumption is small, as the $P_b(s)$ time constant, τ_3, is under five days for the groups. The speed of resumption of smoking does not strictly adhere to an expected relationship with respect to treatment condition: the AC group features the largest τ_3 and the PC group has the smallest. The magnitude of resumption during a quit attempt is quantified with ab, and is relatively small for the groups. Comparing ab and $ab + c'$ values, the mediated pathway and net outcome gains, respectively, it is apparent that the mediated pathway's contribution to the net effect of the quit attempt is consistently small for all of the group averages. Interestingly, the mediated pathways' *relative* contribution to the outcome does not follow a natural progression in terms of relationship to treatment; the mediated pathway's contribution to the outcome is 6.3% for the AC group, 14.9% for ANc, 10.7% for PC, and 13.0% for PNc.

As seen in Fig. 2.6, the AC single subject example appears to successfully quit smoking. The estimated mediation model for this subject is consistent with this observation. Specifically, the magnitude of resumption is near zero (ab equal to -0.24), and the speed at which quit success is achieved is significantly faster than that of the group average counterparts, as the outcome settling time is 10.64 days. Conversely, the PNc single subject example appears to fully relapse (comparatively large b estimate) and does not feature inverse response (positive values for both the a and τ_a values).

The estimated single subject models are generally less accurate, particularly the PNc subject models. This can be attributed to greater variance in the single subject data. The greater degree of variance is evident in Fig. 2.2: both subjects' baseline signals are very noisy, feature *Craving* signals with lower signal-to-noise ratios, and

the PNc subject's resumption also shows greater variance. These data quality issues are typical when considering single subject data and consequently pose a significant challenge to optimization of smoking interventions given that patient-specific models would ideally act as the basis for development of personalized smoking cessation treatments. As parameters in continuous-time models are not a function of the sampling time, a continuous-time system identification approach is appealing in terms of managing data quality issues such as missing data and non-constant measurement intervals—both common characteristics of self-reported behavioral data (McCarthy *et al.*, 2008b; Timms *et al.*, 2012). Discussed in more detail in Chapter 5, future estimation of reliable single subject smoking cessation models may benefit from data collected in clinical trials designed with system identification in mind (Ljung, 1999; Deshpande *et al.*, 2012; Rivera, 2012).

Using an iterative estimation procedure similar to that previously described, models for the "reverse" mediation scenario were also estimated, where *Cigsmked* was the mediator and *Craving* the outcome (models not shown). This similarly resulted in models with high goodness-of-fit values and corresponding simulations that appeared to accurately represent the *Craving* and *Cigsmked* dynamics. This suggests that a *Craving*-*Cigsmked* interrelationship captured in the clinical trial ILD is not fully described by a single mediation model. This significantly motivated identification of an alternative description of the smoking cessation process that accounts for this interrelationship in a parsimonious manner.

Self-Regulation Model

One of the greatest opportunities afforded by ILD is the ability to study self-regulating and self-exciting phenomena (negative and positive feedback, respectively) within behavior change processes (Collins, 2006). Regulatory behaviors have historically been

of significant interest in terms of characterizing addiction generally and smoking behaviors specifically (Carver and Scheier, 1998; Solomon, 1977; Solomon and Corbit, 1974; Velicer *et al.*, 1992). The Nicotine Regulation Model, for example, proposes that smoking is done in an attempt to maintain a smoker's blood nicotine set point, where deviations from this set point are the result of environmental conditions (e.g., cigarettes not readily available; Velicer *et al.* 1992). More complicated but conceptually similar mechanisms propose that cigarettes are smoked in order to regulate emotional states or an $Urge$ set point (Velicer *et al.*, 1992; Solomon, 1977; Solomon and Corbit, 1974; Walls and Rivera, 2009; Timms *et al.*, 2013, 2014c); these emotional states may be directly affected by environmental factors unrelated to the act of smoking (Velicer *et al.*, 1992). However, the ability of such theories to inform effective treatment strategies, such as time-varying adaptive interventions, has been limited. This is largely due to the historically significant challenges associated with effective measurement and mathematical modeling of behavioral dynamics (Riley *et al.*, 2011).

Fig. 2.7 depicts a block diagram of the smoking cessation process that features self-regulation. Generally, Fig. 2.7 suggests that cigarette smoking is done in order to maintain a $Craving$ set point and an attempt to quit smoking is a disturbance on this self-regulatory relationship. Essentially, this block diagram reflects an intuitive process: over time, an increasing desire to smoke leads to smoking activity, which then reduces that desire in the short-term. Specifically, Fig. 2.7 describes a feedback loop in which a biochemical or psychological controller, $C(s)$, responds to the deviation, e, between a craving set point, r_{crav}, and the actual measured craving signal ($e = r_{crav} - Craving$). $Cigsmked$ is a sum of the outputs from $C(s)$, the craving self-regulator, and $P_d(s)$, the effect of the quit attempt; $Cigsmked$ then acts as an input to

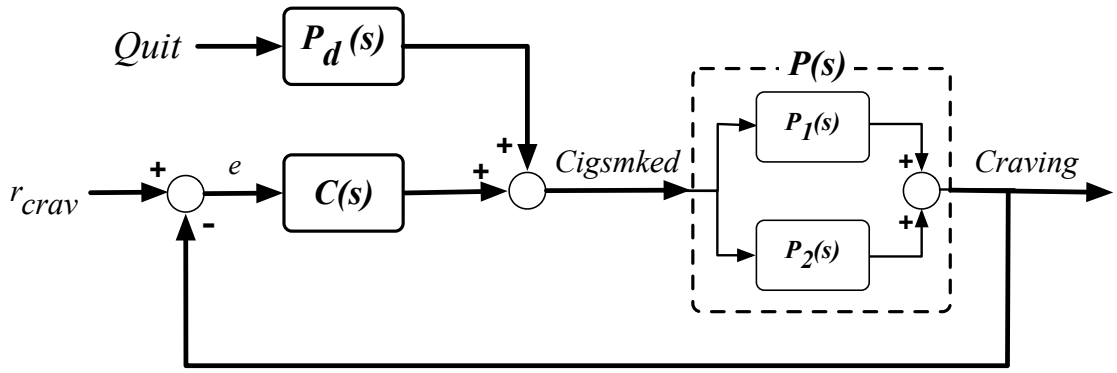

Figure 2.7: Block Diagram Depicting a Smoking Cessation Self-Regulation Model Relating *Craving* and *Cigsmked*.

$P(s)$, producing *Craving*. The associated closed-loop transfer function structures are:

$$Craving = \left(\frac{PC}{1+PC}\right) r_{crav} + \left(\frac{PP_d}{1+PC}\right) Quit \tag{2.28}$$

$$Cigsmked = \left(\frac{C}{1+PC}\right) r_{crav} + \left(\frac{P_d}{1+PC}\right) Quit \tag{2.29}$$

As the output of $P(s)$ is *Craving*, this function will require a system zero, which stems from the sum of two subprocesses in parallel. Fig. 2.7 depicts these underlying, competing processes. Mapping the self-regulatory relationship in Fig. 2.7 to a generalized description of self-regulating behaviors described in Carver and Scheier (1998), $P(s)$ and $C(s)$ correspond to "Effect on Environment" and "Behavior" processes, respectively.

As in the case of the mediation models, estimating the self-regulation models employed a prediction-error approach to obtain continuous-time linear models from sampled data (Ljung, 2009). The system identification procedure is similar to that previously described:

(1) Estimation of $P(s)$ as a single-input single-output (SISO) system with *Cigsmked* as the input and *Craving* as the output according to a given low order transfer function structure.

(2) Simultaneous estimation of $C(s)$ and $P_d(s)$ as a multi-input single-output (MISO) system with $e = r_{crav} - Craving$ and $Quit$ as the inputs and $Cigsmked$ as the output according to given low order transfer function structures for $C(s)$ and $P_d(s)$.

(3) Simulation of the $Craving$ and $Cigsmked$ responses to $Quit$ according to the estimated $P(s)$, $C(s)$ and $P_d(s)$ expressions.

(4) Calculation of $Craving$ and $Cigsmked$ goodness-of-fit according to the criterion in Equation 2.24.

This procedure was implemented in MATLAB through the previously described custom GUI that relied on the `pem` routine, an `idproc` model object, and the process models notation for estimation (Ljung and Singh, 2012). The four steps were repeated for different combinations of transfer function structures, beginning with the lowest order (gain-only) equation structure, with structural complexity increased as necessary. Similarly, various craving set points were examined: baseline $Craving$ level, a linearly decreasing $Craving$ function, and absolute $Craving$ equal to zero. In assessing the group average and single subject candidate model estimates for goodness-of-fit, model parsimony, parameter realizability, and through simulation, the appropriate r_{crav} value was found to equal the baseline $Craving$ level, and the following transfer function structures were found to appropriately represent the observed cessation dynamics:

$$P(s) = \frac{K_1 (\tau_a s + 1)}{(\tau_1 s + 1)} \tag{2.30}$$

$$P_d(s) = K_d \tag{2.31}$$

$$C(s) = \frac{K_c}{(\tau_c s + 1)} \tag{2.32}$$

Table 2.3: Self-Regulation Model Parameter Estimates and Goodness-of-Fit Values.

Treatment	AC	ANc	PC	PNc	AC	PNc
Data Set	Avg	Avg	Avg	Avg	Sgl	Sgl
Craving Fit [%]	87.33	77.65	77.51	62.25	66.90	57.59
Cigsmked Fit [%]	89.16	83.03	91.44	84.12	77.09	62.99
$P(s)$ K_1	0.77	0.74	0.50	0.52	1.57	-2.21
$P(s)$ τ_a	-1.99	-3.76	-14.34	-21.90	-3.05	3.45
$P(s)$ τ_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
$C(s)$ τ_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

The parameter estimates and goodness-of-fit percentages are tabulated for the treatment group averages and the single subjects in Table 2.3. The corresponding model outputs are depicted in Figures 2.5 and 2.6.

As evident in Table 2.3 and Figures 2.5 and 2.6, high goodness-of-fit values and high fidelity simulations are obtained with low order transfer function structures. As before, the negative system zero, τ_a, indicates $P(s)$ represents the sum of two competing subprocesses in parallel. Whereas the group average mediation models suggest the subprocess with the faster speed of response is positive-gained, it is now the negative-gained subprocess that has an immediate speed of response ($P_1(s) = K_{p_1}$, where $K_{p_1} < 0$); conversely, it is the positive-gained function that has the slower speed of response ($P_2(s) = K_{p_2}/(\tau_{p_2}s + 1)$, where $K_{p_2} > 0$ and τ_{p_2} equals τ_1 from Equation 2.30 and Table 2.3). The negative value of K_{p_1} corresponds to the initial increase in *Craving* that results from the quit-induced, step-like initial decrease in *Cigsmked*. The positive value of K_{p_2} corresponds to the settling of *Craving* to

below baseline levels that results from the dramatic reduction in the group average *Cigsmked* signals. As the group average estimates in Table 2.3 for K_1 are positive, it follows that the magnitude of K_{p_2} is greater than that of K_{p_1}. As previously described, the sum of two subprocesses agrees with the concept of dual impulsive and executive neurological processes that compete during a quit attempt (Bickel *et al.*, 2007). Altogether, the craving reduction per unit decrease in daily cigarettes smoked is larger for the active drug groups versus the placebo groups: K_1 equal to 0.77 for the AC group, 0.74 for ANc, 0.50 for PC, and 0.52 for PNc.

The $P_d(s)$ model corresponds to the initial quit-day reduction in *Cigsmked*, and the magnitude of this drop is largest for the AC group (K_d = -15.01) and smallest for the PNc group (K_d = -10.24). As the $C(s)$ function models the post-quit smoking resumption for all of the estimated expressions, it follows that the feedback pathway is responsible for relapse. A distinct advantage of the continuous-time modeling approach is that the nature of the craving self-regulator can be reverse-engineered, as interpretation of the estimated continuous-time $C(s)$ function is straightforward. Evident in Equation 2.32, $C(s)$ is described by a first order differential equation for both the group average and single subject models. Consequently, the craving self-regulator can be classified as a proportional-with-filter controller. This is significant, as such a controller allows offset, and therefore $C(s)$ does not necessarily track the set point (pre-quit baseline *Craving* levels), effectively allowing the model to capture possible quit attempt success as well as failure. The filter component, $1/(\tau_c s + 1)$, serves to attenuate the influence of frequent changes in e on resumption, and suggests that the influence of unmeasured disturbances that lead to the apparent noise in e is suppressed on average. This may support the notion in behavioral science that *Craving* dynamics induced by abstaining from smoking is the dominant factor in determining relapse versus cessation success (McCarthy *et al.*, 2008a; Piper *et al.*, 2008). In gen-

eral, the connection between a fundamental control paradigm (PID control) and a mechanism of change central to behavioral science (self-regulating smoking behavior) makes a case for the potential utility of system identification techniques to behavioral health settings in the future (Ogunnaike and Ray, 1994; Carver and Scheier, 1998; Solomon and Corbit, 1974; Solomon, 1977; Velicer *et al.*, 1992; Timms *et al.*, 2014c). Comparing K_c and τ_c estimates in Table 2.3 for the group average models, the craving self-regulator appears to provide a relatively small and slow contribution to the net change in $Cigsmked$. Interestingly, the overall influence of the feedback pathway on cessation dynamics may be diminished by the combined active drug and counseling treatment, as K_c is approximately 63% smaller for the AC group compared to the ANc and PC groups, and is approximately 73% smaller compared to the PNc group.

The single subject self-regulation models generally have lower goodness-of fit values as compared to the group average models. This was expected, as the ILD from which both mediation and self-regulation models were estimated feature significant variance. (Experimental design options that can mitigate the data quality issue typical of single subject behavioral data is discussed in Chapter 5.) Interestingly, the PNc single subject's $Craving$ and $Cigsmked$ goodness-of-fit values are approximately 10% greater for the self-regulation model compared to the mediation model. This supports the case for self-regulation as a better description of the smoking process.

Focusing on the single subject model estimates specifically, the successful quit attempt for the AC single subject is appropriately represented by the self-regulation model. Although the PNc subject's $P(s)$, $P_d(s)$, and $C(s)$ dynamics are adequately described by the same low order differential equation structures as the group average models, the characteristics of the corresponding parameter estimates contrast those of the group averages. This is due to the subject's failed quit attempt, which is characterized by a full resumption in $Cigsmked$ and a lack of inverse response in

Craving. Whereas the group average model estimates feature K_c estimates that are positive and small in magnitude, the PNc subject's K_c equals -6.25. This large, negative value reflects the significant resumption evident in the subject's *Cigsmked* signal, which is the response to e, the input to $C(s)$ that is essentially the inverted *Craving* signal ($e = r_{crav} - Craving$, where r_{crav} is the baseline *Craving* level). The PNc subject's K_1 estimate is also negative, contrasting the other models' K_1 values; this estimated $P(s)$ model reflects the subject's increase in *Craving* that results from the initial drop in smoking and the subsequent settling to approximately pre-quit levels that results from the apparent relapse. In terms of the underlying subprocesses, where $P(s) = P_1(s) + P_2(s)$ with $P_1(s) = K_{p_1}$ ($K_{p_1} < 0$) and $P_2(s) = K_{p_2}/(\tau_{p_2}s + 1)$ (K_{p_2} also negative and $\tau_{p_2} = \tau_1 = 10.76$ days, per Table 2.3). The relative signs of these subprocess gains are a departure from the relationship inferred for the other models examined.

The feedback nature of the self-regulation model suggests that *Craving* and *Cigsmked* are fundamentally related and a change in one variable results in a change in the other. In other words, the self-regulation models describe a *Craving-Cigsmked* interrelationship that cannot be accounted for by classic mediation. This is significant as a *Craving-Cigsmked* interrelationship was originally suggested by the fact that both the mediation models presented and the reverse mediation models (not shown) have high goodness-of-fit values and high fidelity simulations. Altogether, smoking cessation behavior change is more appropriately and parsimoniously represented as a self-regulatory process as opposed to a mediational relationship. Furthermore, very poor models result from estimation of reverse self-regulation models, where $P(s)$ is said to accept *Craving* as the input, producing *Cigsmked* as the output, etc. This supports the case that Fig. 2.7 is a more appropriate representation of the *Craving* and *Cigsmked* relationship (i.e., *Craving* is the variable being regulated).

Conclusions and Recommendations

Drawing from intensive longitudinal data collected from a clinical trial of bupropion and counseling as aids to smoking cessation, system identification models were developed to describe smoking cessation as mediational and self-regulatory processes. Ultimately, these models differ in how they each describe the resumption of smoking: for mediation, resumption is the result of daily changes in $Craving$ (via $P_b(s)$); for self-regulation, resumption is the result of a craving self-regulator. Analysis of both sets of models highlight the utility of continuous-time system identification in behavioral health settings. Specifically, continuous-time system identification facilitates estimation of parsimonious expressions that accurately represent complex dynamic features within the smoking cessation process. Furthermore, these continuous-time expressions are conducive to straightforward identification and interpretation of the dynamics—in this section, shedding light on the nature of the two competing subprocesses that together form $Craving$, and that the craving self-regulator is a proportional-with-filter controller on average. Altogether, the models developed suggest that self-regulation more appropriately describes the process of smoking cessation. Regardless of structure, parameter values estimated from the group average data, which are signals with high signal-to-noise ratios, suggest both bupropion and counseling have some effect on craving and reduction of smoking behavior.

This section has effectively shown that system identification is useful in analysis of smoking cessation interventions and for comprehensively describing the process of smoking cessation. The dynamical modeling strategy used here could be further applied to the clinical trial data examined here in order to study alternate signal relationships and behavioral mechanisms. Notably, participants in the UW-CTRI clinical trial completed up to seven self-reports at pseudo-random intervals through-

out the day (McCarthy *et al.*, 2008b). These reports assessed environmental factors as well as behavioral states as experienced by the participant at the time the report was completed (whereas the ER focused on a participant's average behavioral state over the previous 24 hours). Behavioral data collected in this fashion is said to reflect the influence of an individual's natural environment, and may measure such influences. Continuous-time models estimated using data from these reports could provide some insight into environmental sources of inter-individual variability, and consequently may act as the basis for more reliable idiographic model estimation. However, the generalizability of such models would still be limited due to the secondary nature of the analysis. In the future, more informative single subject data sets—and ultimately more reliable patient-specific smoking cessation models—could be obtained through a novel smoking cessation clinical trial that draws from experimental design techniques in system identification. Such a trial may vary intervention dosage over time (e.g., bupropion dose, counseling frequency), use self-reports more conducive to measurement of nuanced behavioral dynamics (e.g., *Craving* assessed on a 0 to 100 point scale), or feature a longer self-reporting protocol. However, the medical, practical, and ethical concerns associated with human subjects would also have to be addressed simultaneously. A smoking cessation clinical trial designed to produce more informative single subject data sets would involve experimental design strategies similar to those described in Deshpande *et al.* (2012) and Deshpande and Rivera (2013). Specifically, Deshpande *et al.* (2012) and Deshpande and Rivera (2013) propose optimization-based approaches for the design of periodic, deterministic input signals that facilitate cross-validation, constraint handling, and altogether "patient-friendly" operation.

Ultimately, self-regulatory models similar to those estimated here could inform novel treatment strategies (Timms *et al.*, 2013; Rivera, 2012; Nandola and Rivera,

2013; Riley *et al.*, 2011). Although accurate models of patient-specific behavioral dynamics would ideally be used to personalize interventions, a more practical intervention design approach may rely on a self-regulatory model described in the block diagram in Fig. 2.7, but is parameterized to reflect a representative failed quit attempt—full resumption in *Cigsmked*, inverse-free *Craving*, and no net change in *Cigsmked* or *Craving*. A representative model of a failed quit attempt could be similar to the PNc single subject model examined here. Timms *et al.* (2013) presents such an example of a representative model of quit attempt failure. For designing an adaptive smoking cessation intervention, such a representative model could be used in conjunction with controller design principles to develop an algorithm that defines intervention adjustment (e.g., medication dosage increases), based on a patient's baseline conditions, self-reported smoking and withdrawal symptoms (e.g., daily *Craving* and *Cigsmked* reports), and environmental or other disturbances (e.g., *Stress*). A Hybrid Model Predictive Control approach is appealing as it can simultaneously manage manipulated variables that are on discrete scales (e.g., discrete medication dosage) and constraints (e.g., medication toxicity levels) in an optimal manner (Nandola and Rivera, 2013; Rivera, 2012; Timms *et al.*, 2013). Fig. 2.8 depicts the general form of this control scheme. An optimized, adaptive smoking intervention designed in an HMPC framework could also include features of other well-known control approaches. For example, variables that are effectively non-time-varying but may be relevant to the cessation process, such as the presence of a genetic variant in a patient's nicotine metabolism genes (Chen *et al.*, 2012), could act as scheduling variables in an intervention featuring gain-scheduled Model Predictive Control (Chisci *et al.*, 2003). Finally, event-based control offers controller capabilities that may be appropriate for intervention design. Specifically, event-based control concepts offer a way to mitigate relapse that may otherwise result from time-varying cues to smoke (e.g., proximity

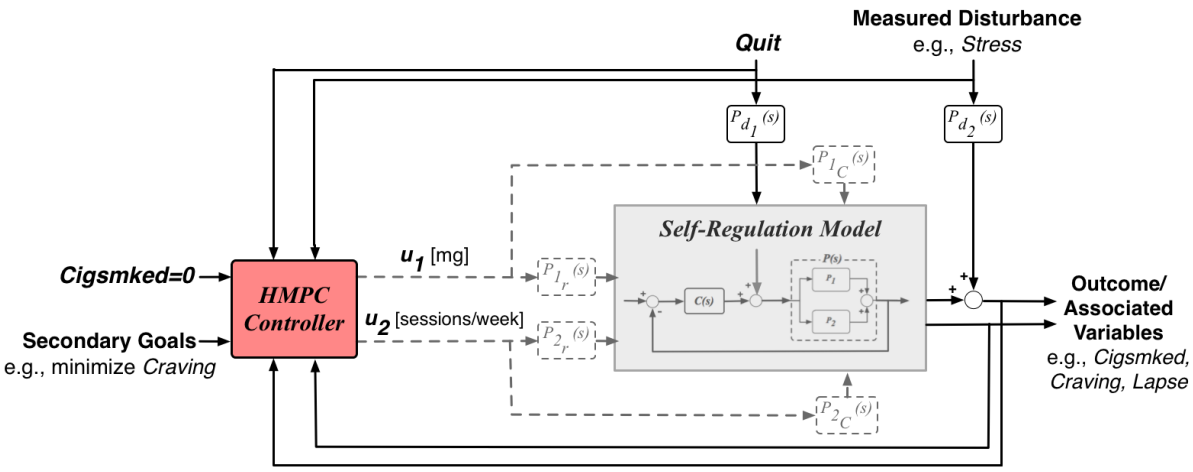

Figure 2.8: Block Diagram of a Hybrid Model Predictive Control Approach to Design of an Optimal, Adaptive Smoking Cessation Intervention.

to smokers) or other disturbances (Pawlowski *et al.*, 2012), and would be particularly appealing for interventions that draw from models of cessation dynamics on shorter time scales (e.g., within-day dynamic models).

2.5 Additional Modeling Considerations

The flexibility of dynamic modeling and system identification techniques suggests they can be useful for investigating a variety of questions in this problem setting. This section outlines some additional opportunities for employing an engineering modeling approach to study smoking behaviors and tobacco interventions.

2.5.1 Characterizing Confounding Influences

The presence of confounding factors are a major concern to behavioral scientists when estimating models of causal mechanisms (Li *et al.*, 2007; Shadish *et al.*, 2002). In dynamical systems terms, confounding variables are exogenous disturbances that may or may not be measured (Timms *et al.*, 2014c,b). Confounders are generally not, or cannot be, controlled for experimentally, and may be static or time-varying.

Even when measured, confounders can increase the complexity of the modeling problem, particularly when considering multivariate causal relationships; for example, Li *et al.* (2007) describes 19 different ways a single confounder can affect a mediational relationship.

The dynamic mediation or self-regulation models developed in the previous section can be adjusted in a straightforward manner in order to characterize the effect of a confounding variable. Consider the confounded mediational relationship depicted in the path diagram in Fig. 2.9; in this example, a confounding variable, Z, influences the levels of both the mediator and outcome variables (Li *et al.*, 2007). The magnitudes of these net effects are represented by ψ and ω, respectively. It can be shown that the structural relationship in Fig. 2.9 can be cast as a dynamical system described by the following:

$$M(s) = P_a(s)X(s) + P_\psi(s)Z(s) + d_1(s) \tag{2.33}$$

$$Y(s) = P_{c'}(s)X(s) + P_b(s)M(s) + P_\omega(s)Z(s) + d_2(s) \tag{2.34}$$

where $P_\psi(s)$ and $P_\omega(s)$ are the transfer functions representing how $Z(s)$ affects $M(s)$ and $Y(s)$, respectively. With ILD for Z, $P_\psi(s)$ and $P_\omega(s)$ can be estimated in addition to $P_a(s)$, $P_b(s)$, and $P_{c'}(s)$. If the resulting set of five estimated functions account for more variance in the observed M and Y signals, as compared to models of the standard three-variable dynamic mediation models (equations 2.22 and 2.23), the case is made that Z is a significant exogenous effects. Ultimately, modeling efforts of this sort are limited by the availability and quality of ILD.

2.5.2 Conceptualizing Mechanisms of Treatment Effects

The group average dynamic self-regulation models presented in Timms *et al.* (2014a) account for the effects of bupropion and counseling within the different param-

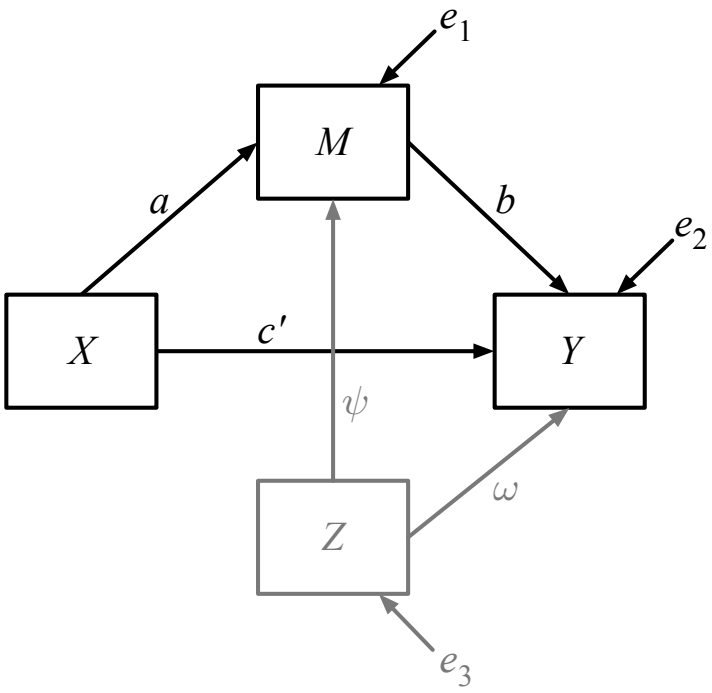

Figure 2.9: Path Diagram Depicting One Manner by which a Single Confounding Variable can Affect a Mediational Relationship (Li *et al.*, 2007).

eter estimates in Table 2.3. Ultimately, dynamic self-regulation models that would be most informative to designing novel treatment regimens would explicitly treat intervention components as independent, exogenous inputs to the psychological feedback system. Future efforts may be able to delineate the treatment effects in this manner. As will be briefly discussed in Chapter 5, ILD from a novel clinical trial designed with system identification principles in mind could facilitate estimation of high fidelity models of this type. However, the models described in previous sections are conducive to general, informal analyses of therapeutic mechanisms.

For example, inspection of Fig. 2.7 suggests there are three clear degrees-of-freedom on which to intervene: r_{crav}, $Craving$, and CPD directly. The seven first-line tobacco dependence medications are all thought to directly reduce $Craving$ (Benowitz, 2009; Tobacco Use and Dependence Guideline Panel, 2008; Lexicomp, 2014; Rennard *et al.*, 2014). Several aspects of counseling likely act most directly

on CPD; for example, counseling may lead a person to reduce their own access to cigarettes (e.g., where they throw away their cigarettes pre-TQD). Bupropion is thought to interfere with nicotine addiction mechanisms (Benowitz, 2009; West *et al.*, 2008); this effect could be modeled as a disturbance on r_{crav}, the psychological $Craving$ set point. As engineers commonly run simulations to explore how to augment a system or alter its operation in order to obtain more favorable outputs, Fig. 2.7 and equations 2.35 and 2.36 can facilitate similar efforts in the context of cessation interventions (Timms *et al.*, 2014c, 2013). Timms *et al.* (2013) provides an example of this opportunity. Here, the self-regulatory process depicted in Fig. 2.7 is considered, but for a hypothetical subject unable to quit smoking on their own. This subject is patterned after the single subject from the PNc group modeled in Timms *et al.* (2014a). Timms *et al.* (2013) then posits that a hypothetical pharmacotherapy acts as a disturbance on the psychological CPD-$Craving$ feedback process featuring dual modes of therapeutic action: (1) directly altering r_{crav} and (2) directly altering $Craving$. Fig. 2.10 depicts this sort of mechanism of effectiveness. The corresponding closed-loop transfer functions are:

$$Craving = \left(\frac{P_{I_c} + P_{I_r}PC}{1 + PC}\right) u_I + \left(\frac{PP_d}{1 + PC}\right) Quit \tag{2.35}$$

$$CPD = \left(\frac{(P_{I_r} - P_{I_c})C}{1 + PC}\right) u_I + \left(\frac{P_d}{1 + PC}\right) Quit \tag{2.36}$$

where u_I denotes whether the hypothetical medication is active ($u_I = 1$) or not ($u_I = 0$), P_{I_c} is the dose-response model representing how u_I affects $Craving$ over time, and P_{I_r} is the dose-response model representing how u_I affects r_{crav} over time. Primarily through simulation, Timms *et al.* (2013) then identifies a combination of first order transfer functions for P_{I_c} and P_{I_r} that together promote successful cessation and decreased $Craving$ values, when u_I is implemented in simulation as a fixed intervention that becomes fully effective on TQD. Simulations of this sort could help

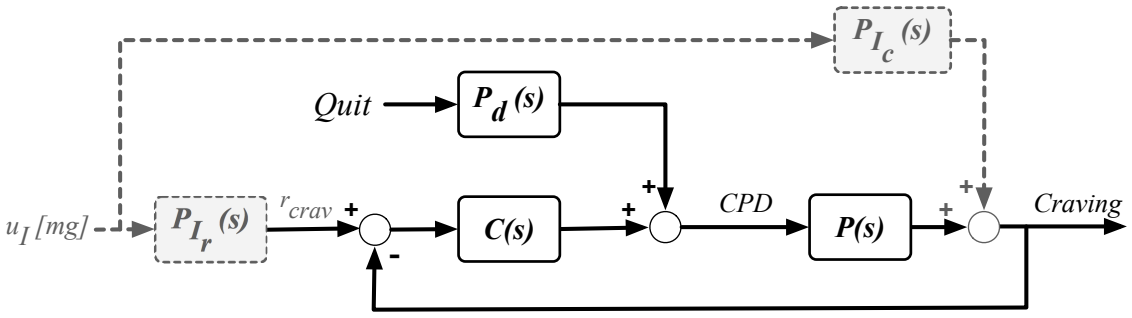

Figure 2.10: Block Diagram Depicting the Mechanism of Action for a Hypothetical Cessation Medication (Timms *et al.*, 2013).

evaluate the most effective mechanisms by which novel therapies should intervene.

2.5.3 Getting a Sense for Parameter Sensitivity

Assuming $r_{crav} = 0$ and $P(s)$, $C(s)$, and $P_d(s)$ adhere to the structures in equations 2.30 through 2.32, the closed-loop transfer functions in equations 2.35 and 2.36 can be rewritten as:

$$Craving = \frac{\left(\frac{K_1 K_d}{1+K_1 K_c}\right)(\tau_a s + 1)(\tau_c s + 1)}{\left(\frac{\tau_1 \tau_c}{1+K_1 K_c}\right) s^2 + \left(\frac{\tau_1 + \tau_c + K_1 K_c \tau_c}{1+K_1 K_c}\right) s + 1} \tag{2.37}$$

$$CPD = \frac{\frac{K_d}{1+K_1 K_c}(\tau_1 s + 1)(\tau_c s + 1)}{\left(\frac{\tau_1 \tau_c}{1+K_1 K_c}\right) s^2 + \left(\frac{\tau_1 + \tau_c + K_1 K_c \tau_c}{1+K_1 K_c}\right) s + 1} \tag{2.38}$$

Clearly, uncertainty in any single or combination of parameters in the $P(s)$, $C(s)$, and $P_d(s)$ transfer functions can have complex effects on $Craving$ and CPD dynamics. This raises a question of how uncertainties in the various parameters will affect the character of the responses. A formal sensitivity analysis lies outside the scope of this dissertation. However, simulation offers a straightforward means for rapidly and informally getting a sense for the effects of small amounts of parameter uncertainty.

For this, the value of each parameter estimate for the PNc single subject self-regulation model (see Table 2.3) was independently increased or decreased by 10%. The responses of $Craving$ and CPD to $Quit$ were then simulated in MATLAB. The

Table 2.4: Percent Change in the Characteristics of the *CPD* and *Craving* Responses to a Step in *Quit* When the Model Parameters are Adjusted by 10%, Relative to the Metrics for the Nominal Case of the Single PNc Subject Estimated Model with No Parameter Uncertainty (See Table 2.3).

	Output:	*Craving*	*Craving*	*Craving*	*Craving*	*CPD*	*CPD*	*CPD*
Parameter	Uncertainty level	Net change	Settling time	Peak magnitude	Peak time	Net change	Settling time	Peak magnitude
$P(s)$, K_1	-10	-2.5	5.6	-7.0	16.1	7.6	5.7	0
$P(s)$, K_1	+10	1.5	-4.9	6.9	-3.0	-6.6	-5.0	0
$P(s)$, τ_1	-10	-1.6	-5.7	4.1	-9.4	-4.3	-5.0	0
$P(s)$, τ_1	+10	1.3	5.1	-3.3	22.9	2.3	4.5	0
$P(s)$, τ_a	-10	1.2	0.7	-0.4	16.1	1.9	0.3	0
$P(s)$, τ_a	+10	-1.7	-0.9	0.9	-3.0	-2.4	-0.4	0
$C(s)$, K_c	-10	8.3	5.6	3.4	16.1	7.6	5.7	0
$C(s)$, K_c	+10	-7.7	-4.9	-2.8	-3.0	-6.6	-5.0	0
$C(s)$, τ_c	-10	0.4	-5.6	-3.1	-4.0	0.9	-5.8	0
$C(s)$, τ_c	+10	-0.6	5.3	3	16.5	-1.2	5.5	0
$P_d(s)$, K_d	-10	-10	0	-10	0	-10	0	-10
$P_d(s)$, K_d	+10	10	0	10	0	10	0	10

net change, settling time, peak magnitude, and time of the peak for both outcome responses were then determined (The MathWorks, 2014h). Table 2.4 documents the percent difference of the responses with parameter uncertainty, relative to the simulated response with no uncertainty.

Per Table 2.4, the time at which the *Craving* peak is reached is generally very sensitive to small perturbations in model parameters in terms of the relative amount of change in that time point. Small amounts of uncertainty in the K_1 and K_c parameters have relatively broad effects, inducing changes of approximately 5% or more in several of the step response characteristics. This agrees with visual inspection of equations 2.37 and 2.38, as the gains, time constants, and damping coefficients in these closed-loop transfer functions are all a function of K_1 and K_c. The net changes in *Craving* and *CPD* induced by a quit attempt, and the peak magnitudes of these signals are also sensitive to small perturbations in K_d. The character of the dynamic responses appear to be relatively robust to uncertainty in τ_a and τ_c. Given that τ_c

corresponds to the frequency of signals the self-regulator filters out, and these simulations do not incorporate a stochastic disturbance, the limited effect of τ_c perturbations is intuitive.

2.5.4 Within-Day Smoking Dynamics

Analysis of the self-regulation process in Timms *et al.* (2014a) and Timms *et al.* (2014c) raises the question, "does a similar phenomenon occur on a within-day level?" Interest in this type of question is ongoing in smoking research settings (Chandra *et al.*, 2011; Shiffman, 2005; Todd, 2004). For example, Chandra *et al.* (2011) characterizes a reciprocal relationship between craving and smoking levels within a day. This is part of a general interest in dynamics on a finer level, where specific determinants of specific instances of smoking could be studied more precisely. However, modeling *within-day* dynamics may incur a number of additional challenges. Intuitively, complex sets of biological, psychological, habitual, environmental, situational, and other types of time-varying disturbances may be more prominent at this level. Consequently, measuring and unentangling these influences may demand significantly more complex models. For example, the effect of stress-management counseling is will be a function of the presence of a stressor at a given time, implying a linear time-invariant model of counseling effects may not be sufficient.

To briefly explore the question of self-regulation on a within-day level, dynamic models were estimated using ILD collected through EMA protocols for approximately 324 subjects in the clinical trial described in Shiffman *et al.* (2006). Although the general relationship depicted in the block diagram in Fig. 2.7 is being considered here, the ILD used in this subsection's modeling work is different from that used in Timms *et al.* (2014a) and Timms *et al.* (2014c), so the craving and smoking signals are represented as *b_cravin* and *cigs* here (per the notation used in alternative statistical

analyses of this data). In deviation variable form, *b_cravin* ranges from approximately -10 to 7; *cigs* was reported for each observation and represents instances of smoking, e.g., a whole cigarette was smoked.

Fig. 2.11 depicts *b_cravin* and *cigs* data for one subject in this study (in deviation variable form). This figure depicts the subject's data as a single time series, where the time scale is minutes since the first self-report was recorded, spanning from minute 0 to minute 79,720. The irregular sampling intervals suggest this modeling effort could benefit from a continuous-time estimation approach (Garnier *et al.*, 2008b; Timms *et al.*, 2014a). However, early efforts indicated that models could not be consistently estimated using routines from the CONtinuous-Time System IDentification (CONTSID) toolbox, a well-established toolbox for estimating continuous-time dynamic models (Garnier *et al.*, 2009, 2008a). Specifically, parameters could only be reliably estimated when a filter was employed with a cutoff frequency that filters out dynamics with time constants around approximately 12 hours.

The limitations of CONTSID and the secondary nature of these analyses motivated estimation of discrete-time models to describe within-day *b_cravin* and *cigs* dynamics. However, discrete-time model estimation routines have difficulty with irregularly sampled data (Garnier *et al.*, 2008b; Ljung, 1999). In order to take advantage of well-established parameter estimation methods and tools such as the System Identification Toolbox (Ljung, 2011), ILD for the approximately 324 subjects is "binned" and averaged. Specifically, a single day is split into seven two hour time intervals, i.e., seven bins: bin 1 corresponds to 7:00 AM to 8:59 AM, bin 2 corresponds to 9:00 AM to 10:59 AM, and so on, through bin 7 that corresponds to 7:00 PM to 8:59 PM. As the clinical trial focused on a six week time period, two weeks pre-quit to four weeks post-quit, each of the 42 days in this time period consists of seven bins. The day relative to TQD was then determined for each observation. Each *b_cravin* and *cigs*

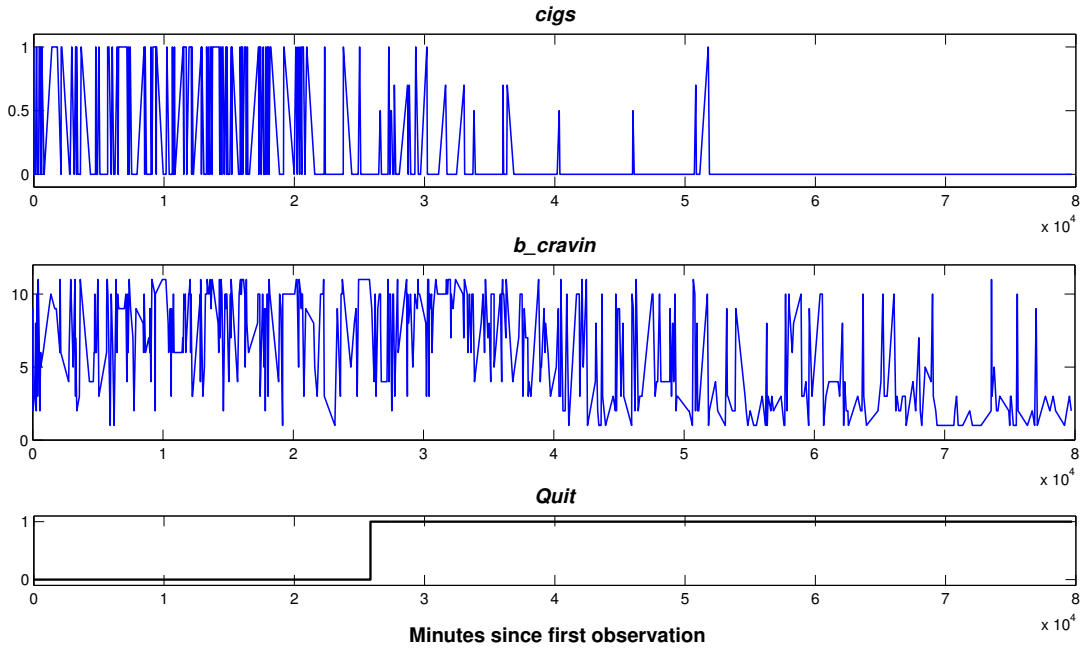

Figure 2.11: Raw *cigs* and *b_cravin* Data for One Subject from the Clinical Trial Described in Shiffman *et al.* (2006).

observation was assigned to the appropriate bin based on the time during the day the report was recorded. The data in each bin for each day relative to TQD was then averaged. The result is 42 sets of time series data, with each time series containing seven data points. Fig. 2.12 depicts binned and averaged time series for several days. As before, there are two identification tasks to be conducted in parallel: (1) estimation of a SISO system where *cigs* is the input and *b_cravin* is the output, and (2) estimation of a two-input single-output (2ISO) problem where *Quit* and $e_{crav} = r_{crav} - b\_cravin$ are the inputs and *cigs* is the output. Here, it is assumed that $r_{crav} = 0$, as was the case for the corresponding day-to-day models. The quit input signal is constant for each time series, equal to 0 for days pre-TQD and 1 for TQD and after. To leverage the parameter estimation methods in MATLAB's System Identification Toolbox, two `iddata` objects are defined for each of the 42 days of data: one corresponding to the SISO problem and a second corresponding to the 2ISO problem. The availability of

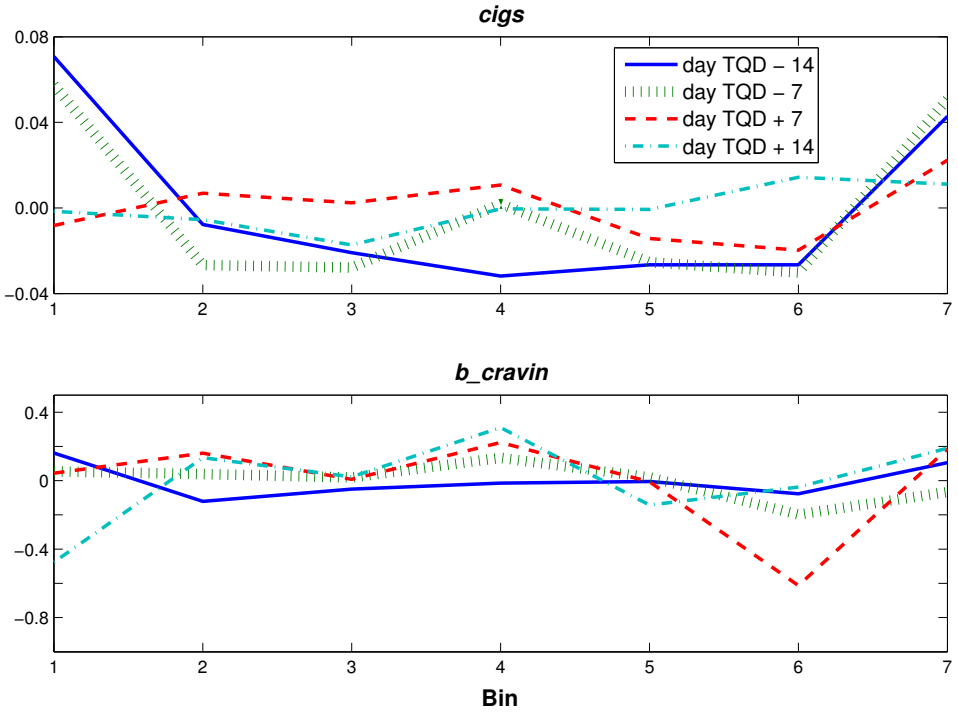

Figure 2.12: Four Days of Binned and Averaged *cigs* and *b_cravin* Time-Series Data from the Study Described in Shiffman *et al.* (2006).

42 time series data sets means there is an opportunity for cross-validation. Here, every other time series is assigned to the estimation data set, and the remaining to the validation data set. Each data set is consequently composed of `iddata` objects for 21 days. For each identification problem, `iddata` objects in the estimation data set are merged into a single object containing 21 "experiments" (The MathWorks, 2014c). This means data for all 21 time series are considered together during the parameter estimation step (Ljung, 2011), as opposed to individual models being estimated for each day.

Discrete-time ARX models in the form of equation 2.2 were then estimated; t in equation 2.2 denotes the current bin. Several combinations of autoregressive and external input polynomial orders were examined (see equations 2.3 and 2.4); model structures with integrators were also considered. Each candidate estimated model

was compared to each of the 21 time series in the validation data set through visual inspection and calculation of a goodness-of-fit percentage.

For the SISO problem, models with fourth order autoregressive polynomials generally did significantly better than models with lower order polynomials. The greatest case was made for an ARX-[4 1 1] model:

$$\begin{aligned} b\_cravin(k) = &-0.09\, b\_cravin(k-1) - 0.5\, b\_cravin(k-2) - 0.3\, b\_cravin(k-3) \\ &- 0.3\, b\_cravin(k-4) - 1.5\, cigs(k-1) + e(k) \end{aligned} \tag{2.39}$$

Of the structures examined, equation 2.39 corresponds to the highest goodness-of-fit value found for a single validation time series, 81%. Furthermore, of the stable models, equation 2.39 corresponds to the highest goodness-of-fit on average, 46%. While the 46% goodness-of-fit average is below the goodness-of-fit values in Table 2.3, this is encouraging as cross-validation is a more rigorous validation method.

Fig. 2.13 features step and impulse responses for equation 2.39, where the respective unit input changes occur at bin 0. The step response indicates that going from 0 cigarettes per bin to 1 per bin for the entire day leads to a relatively large reduction in *b_cravin* initially, but that a unit cigarette smoked during each bin essentially becomes less satisfying as the time since the step change increases. The impulse response suggests smoking a single cigarette quickly leads to a relatively large reduction in *b_cravin*; however, that single cigarette leads to higher *b_cravin* levels four to eight hours later, which seems to follow addiction theory.

The resulting model for the 2ISO problem is:

$$\begin{aligned} cigs(k) = &\ 0.7\, cigs(k-1) + 0.25\, cigs(k-2) - 0.35\, cigs(k-3) + 0.1\, cigs(k-4) \\ &+ 0.3\, cigs(k-5) - 0.02\, e_{crav}(k-1) + 0.02\, e_{crav}(k-2) + e(k) \end{aligned} \tag{2.40}$$

As with the model in equation 2.39, equation 2.40 corresponds to the highest goodness-of-fit on average of the low order structures considered, 51%. It also corresponds to

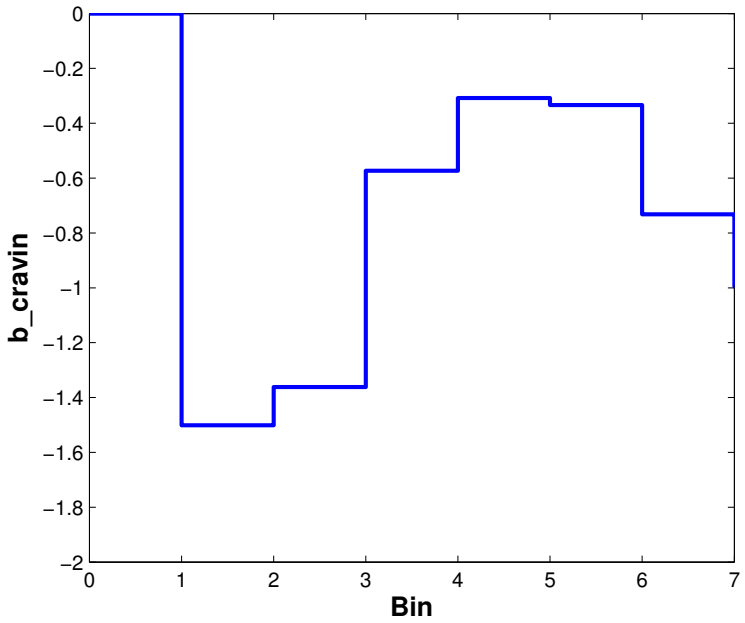

(a) Step Response

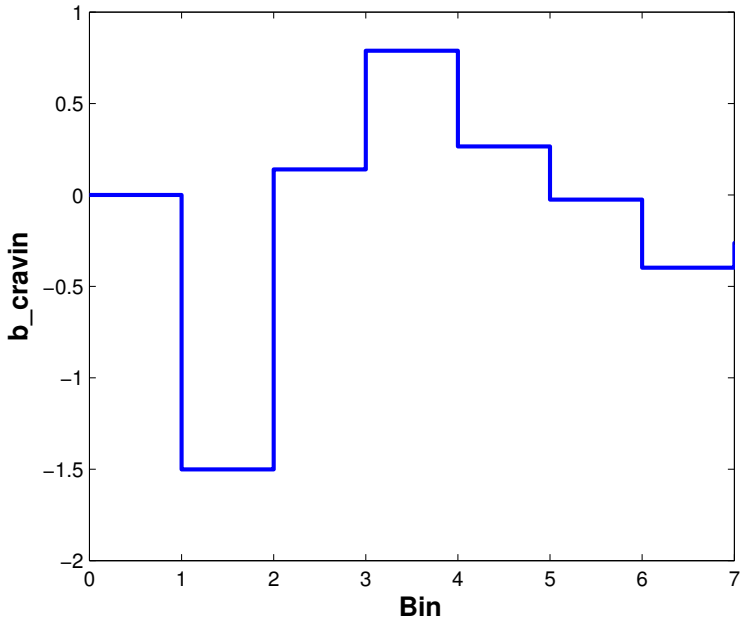

(b) Impulse Iesponse

Figure 2.13: Response of *b_cravin* to a Unit Change in *cigs*.

goodness-of-fits greater than 80% for five of the days in the validation data set, with a 90% fit for one of these days. Note, equation 2.40 attributes no dynamics to the quit path.

The step response in Fig 2.14a suggests that a sustained unit increase in $e_{crav} = r_{crav} - b\_cravin$ leads to a decrease in the number of cigarettes smoked throughout the same day, as would be expected. Fig. 2.14b suggests that a unit decrease in the difference between r_{crav} and b_cravin at time 0 initially leads to a decrease in the amount of smoking, but the overall effect quickly diminishes.

Whereas initiation of the quit attempt acts to excite the CPD and $Craving$ dynamics on a day-to-day scale (see the models in Fig. 2.5 and Fig. 2.6), equation 2.40 suggests initiation of a quit attempt does not excite $cigs$ or b_cravin. The fact that $Quit(k)$ in this model has a constant effect is intuitive though: $Quit$ represents a within-day concept, and conceptually, the transition from not trying to quit smoking to trying to quit occurs sometime between the end of bin 7 on day =TQD-1 and the beginning of bin 1 on day =TQD. However, one might have expected that $Quit$ contributes to equation 2.40 in a way that scales $cigs(k)$ down from pre-TQD baseline levels of $cigs$ to approximately 0 $cigs$, as $P_d(s)$ in equation 2.31 served to do. Future model estimation efforts drawing from novel clinical trial data may provide greater insight into this unintuitive result. Furthermore, such future experimentation amy be able to shed light on the excitation source that induces $cigs$ and b_cravin dynamics, as seen in Fig. 2.12. It has been documented that there is a natural evolution of craving and smoking levels throughout the day in smokers not trying to quit (Chandra *et al.*, 2011); this may be an appropriate starting point for future experimentation and modeling efforts. Such efforts to essentially estimate $cigs(z^{-1})/u_{day}(z^{-1})$ and $b\_cravin(z^{-1})/u_{day}(z^{-1})$—where $u_{day}(z^{-1})$ is an impulse defining the beginning of a new day—were briefly explored using data from the McCarthy *et al.* (2008b)

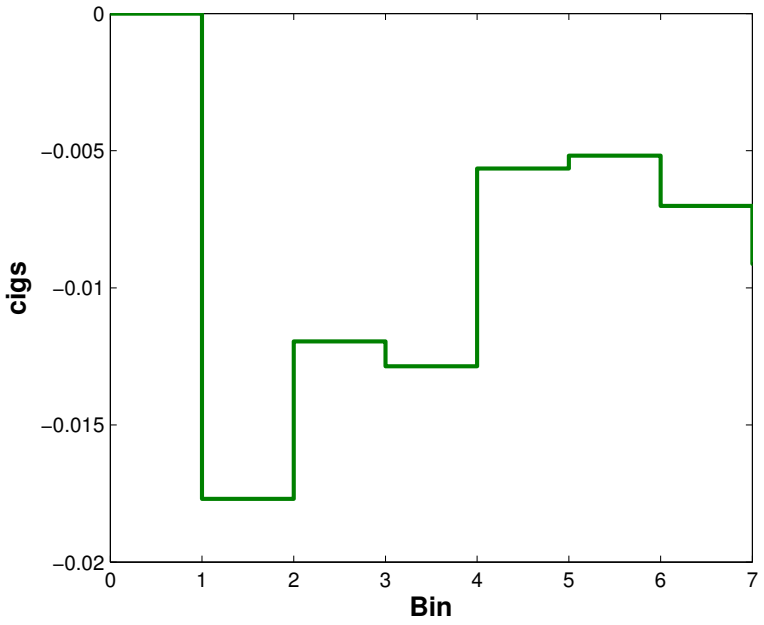

(a) Step Response

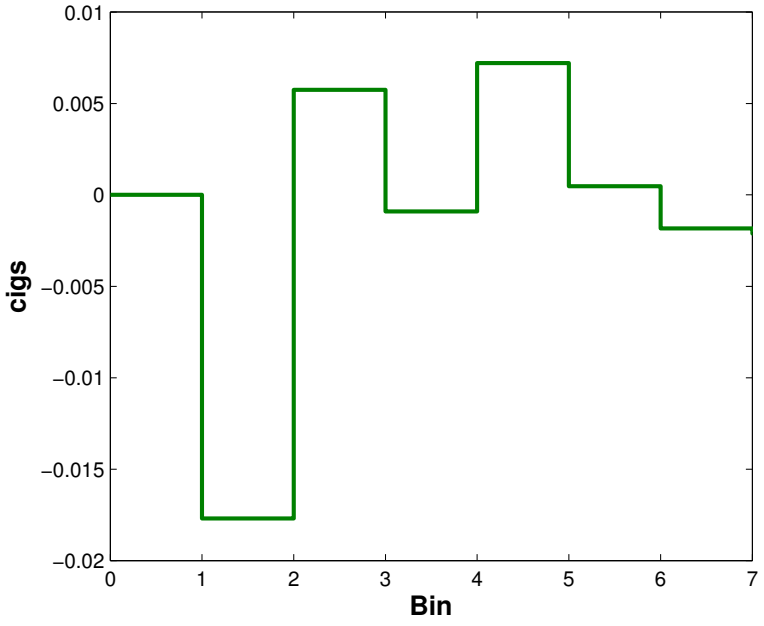

(b) Impulse Response

Figure 2.14: Response of *cigs* to a Unit Change in e_{crav}.

study. However, early analyses indicated that data quality issues prohibit estimation of meaningful models of this sort.

Chapter 3

FORMULATION OF AN HMPC-BASED SMOKING CESSATION INTERVENTION

3.1 Overview

Control systems engineering is playing an increasing role in clinical environments with the goal of automating high quality care, increasing the efficiency of treatments, or treating patients in remote areas or developing nations through telehealth technologies (Doyle *et al.*, 2011). Fraser Health (2014) and Soltesz (2013) use a PID control algorithm to automate delivery of anesthesia; Denai *et al.* (2009) and Ross *et al.* (2009) employ a fuzzy logic controller to automate delivery of medication to cardiac intensive care patients; a predictive control approach to designing treatment regimens for pharmacological HIV therapies is presented in Zurakowski and Teel (2006); Doyle *et al.* (2014) surveys the current landscape of control approaches for use within artificial pancreas technology; and Brier *et al.* (2010) and Gaweda *et al.* (2008) take a Model Predictive Control (MPC) approach to treating anemia.

Recently, predictive control concepts have been introduced into behavioral health settings by casting the task of developing adaptive, behavioral intervention decision rules as a control systems engineering problem (Deshpande *et al.*, 2014, in press; Dong *et al.*, 2012, 2013, 2014; Noble, 2014; Rivera, 2012; Rivera *et al.*, 2007; Savage *et al.*, in press; Timms *et al.*, 2014d). Accomplishing this draws from a number of clear parallels between adaptive behavioral interventions and control systems engineering principles. Table 3.1 summarizes these conceptual connections.

Briefly described in Section 1.4.3, decision systems based on MPC and Hybrid

Table 3.1: Conceptual Connections Between Adaptive Interventions in Behavioral Health and Control Systems Engineering Principles.

Adaptive behavioral intervention component	**Control systems engineering equivalent**
Intervention outcomes	Controlled variables, associated variables
Intervention goals	Set points, controlled variable targets
Tailoring variables	Measured outputs, disturbances, and scheduling parameters (feedback and feedforward signals)
Treatment components (e.g., medication)	Manipulated variables
Clinical use guidelines and restrictions	Constraints
Decision rules	Controller
Behavior change process	Open-loop dynamical systems models

MPC (HMPC) offer distinct advantages in terms of frameworks for adaptive behavioral interventions. Some of these advantageous features include the following: control action is determined by minimization of an objective function (Camacho and Bordons, 1999; Goodwin *et al.*, 2005; Ogunnaike and Ray, 1994; Rossiter, 2003), which is conducive to optimized, personalized dosing (Nandola and Rivera, 2013; Rivera *et al.*, 2007; Timms *et al.*, 2014d); HMPC specifically can manage the pre-determined and discrete-valued nature of treatment component dosages when determining optimal dosing regimens (Bemporad and Morari, 1999; Deshpande *et al.*, 2014; Dong *et al.*, 2013; Nandola and Rivera, 2013; Timms *et al.*, 2014d); the predictive and receding-horizon nature of MPC means dosing decisions can be made with a patient's past, present, and future needs in mind (Deshpande *et al.*, 2014; Dong *et al.*, 2013; Nandola and Rivera, 2013; Timms *et al.*, 2014d), which facilitates "just-in-time" decision making (Intille *et al.*, 2003; Kumar *et al.*, 2013; Riley *et al.*, 2011; Timms *et al.*, 2014d); predictive control can systematically manage multi-input, multi-output (MIMO) systems; and constraints are explicitly considered at each decision point (Camacho and

Bordons, 1999; Goodwin *et al.*, 2005; Nandola and Rivera, 2013; Rossiter, 2003).

In this chapter, an HMPC-based adaptive smoking cessation intervention algorithm is designed with three general tasks in mind:

- **Tracking intervention targets:** The controller assigns adjustments to treatment component dosages in order to take measured treatment outcomes toward target levels.

- **Managing measured disturbances:** The intervention controller seeks to reject measured disturbances that may otherwise move the measured outcomes away from target levels or, more generally, negatively influence performance.

- **Rejection of unmeasured disturbances:** The intervention controller will manipulate treatment component dosages in order to manage the risks posed by unmeasured factors, such as exogenous unmeasured and unmodeled disturbances and plant-model mismatch (which are managed by the controller in a similar manner). The ability of the controller to mitigate the potentially negative effects of such disturbances is critical, given the significant patient-to-patient variability present in this problem setting.

More specifically, this chapter lays the conceptual and computational groundwork for a control engineering-based decision system that personalizes treatment over time in order to meet the changing needs of a smoker trying to quit. An HMPC-based algorithm is derived where CPD and $Craving$ are the primary controlled variables and $Quit$ and $Stress$ are measured and anticipated disturbances. The HMPC algorithm primarily assigns day-to-day adjustments to dosages of treatment components that are both pharmacological and behavioral in nature. Altogether, the MIMO controller formulation is developed to take CPD and $Craving$ to target levels of zero each day.

It should be noted that Timms *et al.* (2014d) proposed an initial framework for an engineering-based smoking cessation intervention algorithm. The decision framework developed in this chapter significantly builds upon, and refines the basic ideas discussed in Timms *et al.* (2014d). Chapter 4 similarly provides a much more thorough analysis of intervention performance and a study of intervention robustness.

This chapter is organized as follows: Section 3.2 outlines the general requirements, components, and structure of the controller designed in this dissertation; Section 3.3 offers details on the open-loop dynamical systems models that act as the basis for design of the controller; Section 3.4 details the decision-making procedure and computations; and Section 3.5 extends the fundamental HMPC framework to one that features three-degree-of-freedom (3DoF) tuning capabilities.

3.2 General Structure, Components, and Requirements of the Intervention

This controller design problem must address a number of factors not seen in more conventional engineering settings. This is largely due to the clinical setting in which the controller will be implemented and the fact that the "plant" in the control loop is a smoker trying to quit. In terms of control-based behavioral health interventions in general, determination, implementation, and ultimately clinical validation of the control actions defined by the controller must adhere to an array of medical, practical (largely logistical), and ethical restrictions. However, smoking interventions involve a number of controller design considerations not present within other behavioral health problem settings (Deshpande *et al.*, 2012; Dong *et al.*, 2012, 2013; Rivera *et al.*, 2007; Rivera, 2012). These factors include the fact that the cessation intervention revolves around a central event—initiation of the quit attempt on a pre-determined target quit date—in which the goal is to support the transition of a behavior from fully "on" to fully "off". This quit process moves CPD toward the cessation goal initially while

simultaneously moving *Craving* further from the intervention goals (Timms *et al.*, 2014a,c). Furthermore, initiation of a quit attempt acts as an exogenous disturbance on the system and also directly corresponds to set point changes on TQD; both of which can be anticipated during every control decision, even for weeks prior to TQD. Additionally, the set points are equal to physically realizable limits on the system (e.g., $CPD = 0$) in this control problem, which may have implications for feasibility of the optimization computations and dosing decision-making.

Altogether, the intervention algorithm derived in this chapter considers the clinical requirements, concerns, and restrictions developed in the following subsection.

3.2.1 General Requirements

Table 3.1 describes the general parallels between an adaptive behavioral intervention and a control loop with combined feedback-feedforward action. Translation of the adaptive smoking cessation intervention concept into a control systems engineering setting must consider the following general clinical and practical aspects of the cessation problem.

- The primary goal of smoking cessation intervention is to facilitate a patient's pursuit of a successful quit attempt. In order to directly support abstinence, dosages should consider CPD (number of cigarettes smoked per day) as a controlled variable. While quit attempt success does not necessarily require complete abstinence for all times beginning on TQD, the primary intervention target is $CPD = 0$, $t \geq$ TQD.

- Although the most straightforward metric of quit attempt success is CPD, the intervention should also seek to mitigate risk factors that could otherwise promote relapse. Notably, high craving and withdrawal levels that are sustained

for long periods or experienced as acute episodic events during a quit attempt are associated with greater rates of lapse and relapse (Allen *et al.*, 2008; Baker *et al.*, 2012; Shiffman *et al.*, 2006). Furthermore, recent work described previously and in Timms *et al.* (2014a,c) notes an interrelationship between $Craving$ (average daily craving levels) and CPD during an attempt to quit smoking. Consequently, $Craving$ is a second controlled variable here. A secondary goal of reducing $Craving$, specifically moving $Craving = 0$, $t \geq$ TQD, is therefore also considered.

- Many smokers attribute lapse events to stress (Marlatt and Gordon, 1980; Shiffman and Waters, 2004; Shiffman, 1982), and rapid changes in stress levels are associated with lapse and relapse (Shiffman and Waters, 2004). Consequently, stress levels should be incorporated dosing decision-making.

- In pursuit of an intervention that is optimized in a personalized manner, assessments of a patient's changing needs should act as feedback and feedforward signals. Specifically, CPD and $Craving$ are the controlled variables (measurements of which act as feedback signals) and $Stress$ is a disturbance (measurements of which act as a feedforward signal). These signals are considered to be self-reported measurements assessed via a mobile phone application or similar mHealth technology.

- It is common for TQD to be determined or assigned days to weeks in advance (McCarthy *et al.*, 2008b; Lexicomp, 2014). Given the significant influence initiation of a quit attempt has on CPD and $Craving$, and its ability to be appropriately represented as an exogenous disturbance (Timms *et al.*, 2014a,c), $Quit$ should be incorporated into the control architecture as a measured and anticipated feedforward signal.

- Given the expense of treatment components, potential side-effects of pharmacotherapies (Tobacco Use and Dependence Guideline Panel, 2008; Lexicomp, 2014; Rennard *et al.*, 2014), and a concern that excessive treatment beyond that required for cessation success may ultimately prove counterproductive (especially in patients who are able to achieve relatively successful outcomes without the aid of an intervention; Rivera *et al.*, 2007), an intervention that could explicitly address a concern for the total dosing demands over a given time period would be desirable.

- A "just-in-time" intervention describes treatment in which anticipatory dosing is implemented *prior* to a potential detrimental change in a patient's needs (Kumar *et al.*, 2013; Timms *et al.*, 2014d).

- The first line medications used for smoking cessation are only available and prescribed in pre-determined, discrete dosages (Tobacco Use and Dependence Guideline Panel, 2008; U.S. Department of Health and Human Services and Centers for Disease Control and Prevention, 2014; Lexicomp, 2014; Rennard *et al.*, 2014). Similarly, cessation treatment components that are behavioral or cognitive in nature can realistically only be delivered in discrete-valued dosages (e.g., the patient does or does not participate in a counseling session in a given day). A clinically-relevant predictive controller should be formulated for a hybrid linear dynamical system (Bemporad and Morari, 1999; Nandola and Rivera, 2013).

- Combination therapies generally have the most success in supporting smoking abstinence (Tobacco Use and Dependence Guideline Panel, 2008; Piper *et al.*, 2009; U.S. Department of Health and Human Services and Centers for Disease Control and Prevention, 2014). Consequently, an intervention formulation

should involve a multi-input system. This, combined with the dual objectives of minimizing CPD and $Craving$, implies that the intervention formulation must be designed around a MIMO system. Given its ability to account for multiple control objectives and manage multiple manipulated variables, an intervention employing an MPC structure should be able to manage the MIMO nature of the system in a systematic manner (Camacho and Bordons, 1999; Goodwin *et al.*, 2005; Rossiter, 2003; Ogunnaike and Ray, 1994).

- Because of the clinical and human health context, "hard" constraints should be explicitly incorporated into determination of the control action. This intervention requirement furthers the case for MPC as the algorithmic framework (Camacho and Bordons, 1999; Goodwin *et al.*, 2005; Rossiter, 2003; Ogunnaike and Ray, 1994).

- Ideally, the dosing schedules assigned by the HMPC-based intervention should not deviate too significantly from current protocols associated with specific treatments. Clinicians may be uncomfortable with such significant deviations from current clinical practice and large degrees of "off label" medication use, likely resulting in challenges in terms of practitioner buy-in.

- The intervention algorithms employed within a clinical setting should feature "clinician-friendly" tuning. As those implementing the intervention will be healthcare practitioners as opposed to engineers, a tuning strategy should be defined that focuses on a subset of tuning knobs and consists of straightforward tuning heuristics.

While the points made above are not necessarily an exhaustive list of the requirements and ideal features of a clinically-implementable adaptive smoking cessation intervention, they reflect the major factors to be considered. Consequently, conceptual and

computational development of an HMPC-based intervention that reflects these requirements and concerns will offer an advanced starting point from which efforts to experimentally assess the clinical utility of this approach can begin.

3.2.2 Treatment Components

The intervention components considered in this work are as follows:

1. **Counseling,** u_c: Brief tobacco dependence treatment delivered in the form of telephone counseling has been shown to be effective (U.S. Department of Health and Human Services and Centers for Disease Control and Prevention, 2014). This treatment component would entail "practical counseling," which includes identification of a patient's "danger situations" that may promote relapse, and development of coping skills (Tobacco Use and Dependence Guideline Panel, 2008). The prominence of telephone quit lines within public health programs (Centers for Disease Control and Prevention, 2014; U.S. Department of Health and Human Services and Centers for Disease Control and Prevention, 2014) suggests this may be an appealing venue through which such counseling could be implemented. Recent efforts to allow an exchange of patient information between quit lines and the servers on which electronic medical records are held (Adsit *et al.* (2014); on which an intervention algorithm may also calculate dosing decisions) may even allow quit line counselors to call patients for whom the HMPC algorithm assigns counseling for a given day.

2. **Bupropion,** u_b: One of the two non-nicotinic first line medications, sustained-release bupropion has consistently been shown to help people quit smoking (U.S. Department of Health and Human Services and Centers for Disease Control and Prevention, 2014). Current bupropion dosing protocols consist of 150 mg once

daily for three days, then increased to twice daily, with eight hours in between doses. Evidence-based bupropion treatment regimens consist of 7-12 weeks of 300 mg per day, with the protocol beginning at least one week pre-TQD (Tobacco Use and Dependence Guideline Panel, 2008; Lexicomp, 2014).

3. **Lozenges,** u_l: One of the five nicotine-based first line cessation medications (Tobacco Use and Dependence Guideline Panel, 2008), nicotine replacement lozenges deliver nicotine orally. Patients assessed to have a high nicotine dependency (e.g., those who smoke within 30 minutes of waking) are directed to take 4 mg lozenges as opposed to 2 mg lozenges. The typical dosing schedule involves one lozenge every one to two hours, up to five lozenges every six hours or 20 per day; one lozenge dissolves in the mouth in approximately 30 minutes (Rennard *et al.*, 2014). Patients typically report that lozenges are less satisfying than a cigarette. This is likely a reflection of the fact that nicotine delivery via the blood stream (initially through tissues in the mouth) occurs on a much longer time scale than nicotine delivered through smoking (Piper *et al.*, 2009).

In summary, the HMPC algorithm must assign the manipulated variables according to the following discrete-valued levels:

$$u_c(k) \in \{0, 1\} \text{ [sessions/day]} \tag{3.1}$$

$$u_b(k) \in \{0, 1, 2\} \text{ [150 mg doses/day]} \tag{3.2}$$

$$u_l(k) \in \{0, 1, 2, ..., 20\} \text{ [lozenges/day]} \tag{3.3}$$

3.2.3 Constraints

As the system being intervened upon is human health, dosing decisions need to adhere to a variety of constraints. These constraints reflect medical concerns (e.g., medication toxicity levels), resource-use limitations (e.g., restrictions on counseling

availability as imposed by a patient's health insurance policy), and practical and logistical considerations (e.g., manufacturer-determined bupropion dosage levels). The intervention developed here considers the following constraints:

- Treatment component dosages can only be assigned in discrete levels, which are pre-determined: u_c can only be assigned as whole counseling sessions; bupropion is only widely available in 150 mg doses (Lexicomp, 2014); and u_l can only be assigned as whole lozenges.

- A minimum number of zero doses of any of the treatment components can be assigned each day. For practical and logistical reasons, a maximum of one dose of counseling can be assigned each day. Dictated by medication safety concerns, a maximum of two 150 mg bupropion doses can be assigned each day (Lexicomp, 2014). To avoid nicotine toxicity symptoms, a patient can be assigned a maximum of 20 lozenges per day (Tobacco Use and Dependence Guideline Panel, 2008; Rennard *et al.*, 2014).

- There is no limit to how much u_c or u_l can increase or decrease from one day to another, other than the limits due to the availability of discrete-valued dosage levels.

- Although up to two doses of 150 mg of bupropion can be assigned to a patient per day, bupropion use protocols require that a patient take one 150 mg dose per day for three days before a second daily dose can be assigned (Tobacco Use and Dependence Guideline Panel, 2008; Lexicomp, 2014). There are no medically-necessitated restrictions on the magnitude at which u_b can decrease from one day to another.

- Assigning unlimited counseling for the duration of an intervention is impractical.

Similarly, a health insurance company will often cover only a certain number of counseling sessions over a given quit attempt (Centers for Disease Control and Prevention, 2006). Because of this, a maximum bound on the total amount of counseling received by a patient over the intervention time period considered should be enforced.

- Negative CPD and $Craving$ values are not physically realizable or logical, respectively. Therefore, the controller should adhere to a lower bound of zero for both controlled variables.

- The results of the clinical trial documented in Piper *et al.* (2009) suggest a synergistic effect when nicotine replacement therapy is combined with bupropion. To maximize the effect of both and reasonably assume a linear system despite a potential non-linearity introduced by the interaction between these treatment components, lozenges should only be assigned when bupropion is also being assigned.

As will be shown in Section 3.4, these requirements will be incorporated into the models representing the open-loop cessation process as a hybrid dynamical system, and as upper and lower limits on the controlled variables, manipulated variables, and move sizes.

3.2.4 Intervention Structure & Decision-Making Process

The closed-loop intervention is patterned after the conceptual structure depicted in Fig. 3.1. Such an approach is particularly appealing given the emergence of technologies that offer cost-effective platforms through which information on changing patient needs can be collected, a decision support system may be operationalized, and dosing decisions may be delivered (Adsit *et al.*, 2014; Aveyard and Raw, 2012;

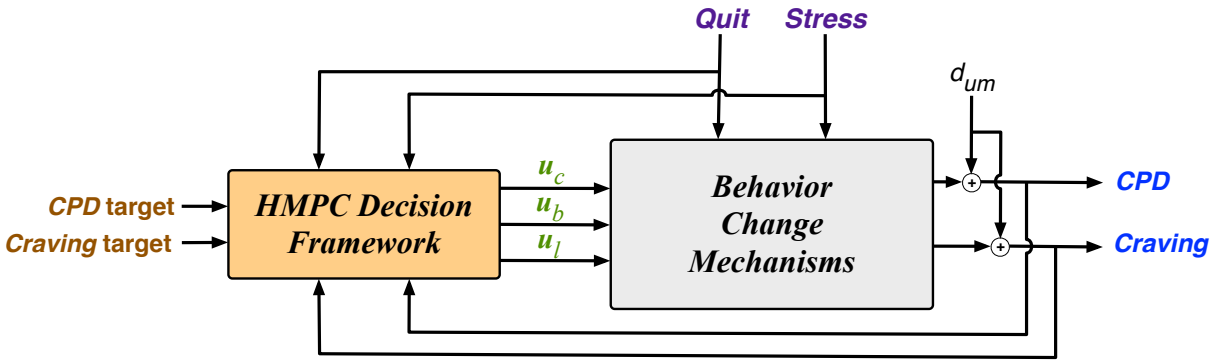

Figure 3.1: Block Diagram of the General Decision Framework for an Adaptive, Smoking Cessation Intervention that Employs an HMPC Algorithmic Structure to Define Daily Dosages of Counseling, Bupropion, and Nicotine Replacement Lozenges.

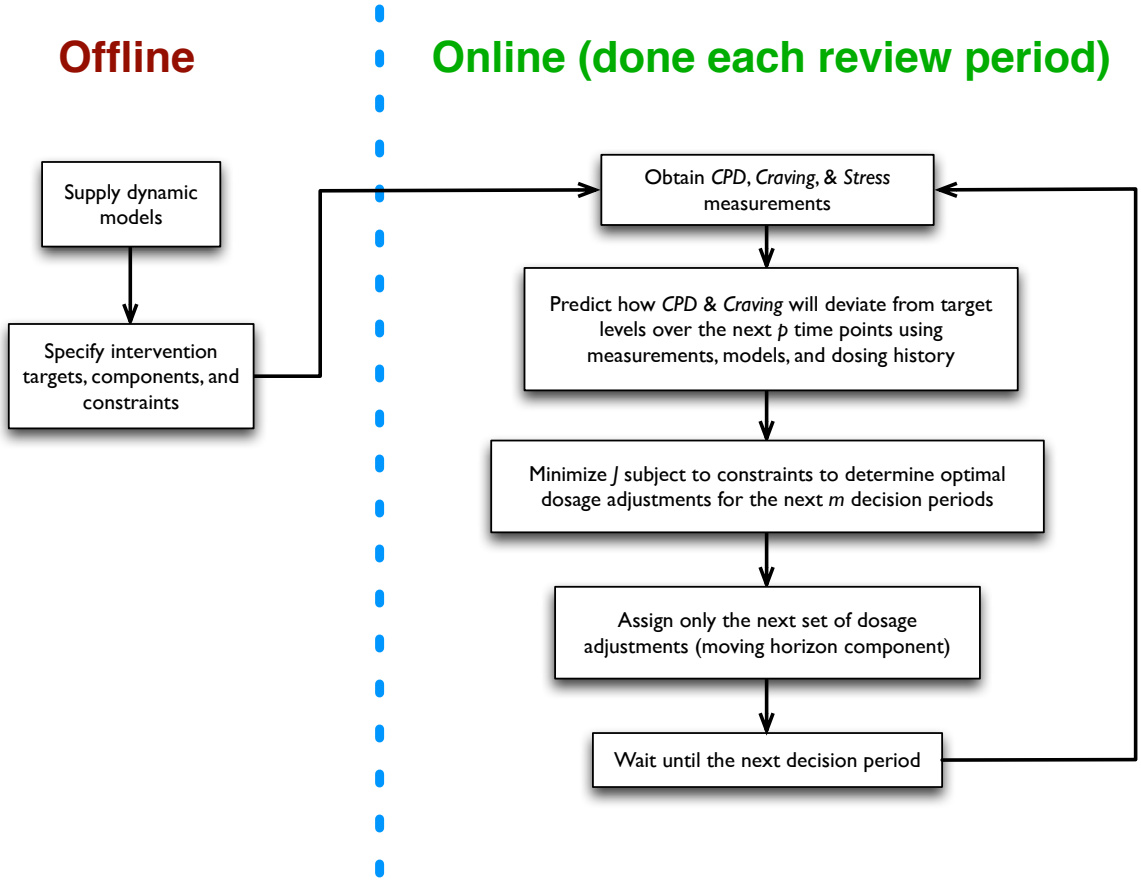

Figure 3.2: Flow Chart of the Decision-Making Process for the HMPC-Based Smoking Cessation Intervention. p is the Prediction Horizon, m is the Move Horizon, and J Denotes the Objective Function.

Kumar *et al.*, 2013; Riley *et al.*, 2011; Tobacco Use and Dependence Guideline Panel, 2008). As alluded to previously, the intervention pursued here focuses on a daily time scale and relies on an HMPC framework.

The decision-making process consists of the steps depicted in the flow chart in Fig. 3.2. The open-loop behavior change models and intervention properties are specified prior to any decision-making. At each decision period, values of the controlled (CPD, $Craving$) and measured disturbance variables ($Quit$, $Stress$) are obtained through patient self-reports via a mobile phone application. The next p days (i.e., the length of the prediction horizon) of CPD and $Craving$ offsets in the absence of additional control action are predicted based on the measurements, nominal models, and previous dosage assignments. The next m days of u_c, u_b, and u_l values are determined by minimizing J subject to constraints (where m is the move horizon and J is the objective function). Only the first set of these treatment dosage levels are assigned, before these steps are repeated on the next decision period. Details of the HMPC formulation are described in the following subsections.

3.3 Open-Loop Dynamical Systems Models

3.3.1 Representative Patient Dynamics

Development of a model-based control algorithm requires a sufficient understanding of the relationship between controlled variables and both manipulated and non-trivial disturbance variables. In this work, the HMPC decision framework that assigns daily dosage adjustments considers the following general transfer function represen-

tation of the system:

$$\begin{bmatrix} CPD \\ Craving \end{bmatrix} = \begin{bmatrix} P_{cpd_c}(s) & P_{cpd_b}(s) & P_{cpd_l}(s) \\ P_{crav_c}(s) & P_{crav_b}(s) & P_{crav_l}(s) \end{bmatrix} \begin{bmatrix} u_c \\ u_b \\ u_l \end{bmatrix} + \begin{bmatrix} P_{cpd_Q}(s) & P_{cpd_S}(s) \\ P_{crav_Q}(s) & P_{crav_S}(s) \end{bmatrix} \begin{bmatrix} Quit \\ Stress \end{bmatrix} \tag{3.4}$$

where $P_{cpd_c}(s)$, $P_{cpd_b}(s)$, and $P_{cpd_l}(s)$ represent the open-loop transfer functions describing how CPD responds to unit changes in the respective manipulated variable dosages. $P_{crav_c}(s)$, $P_{crav_b}(s)$, and $P_{crav_l}(s)$ represent the open-loop transfer functions describing how the patient-reported average daily craving level responds over time to unit changes in the manipulated variables. $Stress$ represents patient-reported levels of average stress per day, which acts as a disturbance to the system. $P_{cpd_S}(s)$ and $P_{crav_S}(s)$ are the corresponding disturbance models. $Quit$ is a signal representing the transition from not attempting to quit smoking to attempting to quit. Here, the $Quit$ signal is treated as the primary measured disturbance, the entire character of which is known prior to any dosing. $P_{cpd_Q}(s)$ and $P_{crav_Q}(s)$ are the corresponding transfer functions representing how a patient's CPD and $Craving$ levels would vary from day-to-day when trying to quit smoking without the aid of any treatments.

Ideally, a patient-specific HMPC-based cessation intervention would be formulated and implemented for each individual patient using open-loop models representative of that smoker's actual individual dynamics. However, estimating these models involves significant experimental and logistical challenges (Deshpande, 2014; Deshpande and Rivera, 2013). Instead, the smoking intervention is developed using nominal models of a hypothetical, representative patient. This representative patient draws from dynamics observed for one subject in the McCarthy *et al.* (2008b) clinical trial and

modeled in Timms *et al.* (2014a). Prior to TQD, this hypothetical subject smokes 9.25 cigarettes per day and experiences an average daily craving level of 16.40. The quit attempt dynamics, which describe an unsuccessful quit attempt, are represented by the continuous-time transfer function structure in equation 3.5 and the corresponding parameters in Table 3.2:

$$P_{y_Q}(s) = \frac{K_Q(\tau_{a_{1_Q}} s + 1)(\tau_{a_{2_Q}} s + 1)}{\tau_Q^2 s^2 + 2\tau_Q \zeta_Q s + 1} \tag{3.5}$$

where y will be either CPD or $Craving$. These models correspond to simplified and rearranged versions of equations 2.37 and 2.38 for the hypothetical patient.

This subject is initially able to quit smoking on TQD without the aid of any intervention, but this also corresponds to an initial increase in $Craving$. However, the patient gradually resumes smoking over time with a corresponding to a reduction in $Craving$ levels. Both outcomes ultimately settle to approximately pre-TQD levels. These open-loop models reflect how initiation of a quit attempt initially brings CPD levels toward the target of cessation, while conversely $Craving$ is initially pushed away from the target of reducing $Craving$.

The open-loop dose-response models for the hypothetical subject are represented as low order continuous-time functions described below. These open-loop dose-response and $Stress$-response models should ultimately be obtained by estimating and validating models using data from novel clinical trials designed with system identification in mind. However, conducting trials of this nature would require significant effort, time, and funding. These considerations, and limitations associated with secondary analysis of the previously-described UW clinical trial (McCarthy *et al.*, 2008b), mean that the following dose-response and $Stress$-response models are primarily informed by literature and analysis of step and impulse responses.

The continuous-time models representing how one 10-minute phone counseling

session leads to changes in either controlled variable employ the transfer function structure in equation 3.6 and the corresponding parameters in Table 3.2:

$$P_{y_c}(s) = \frac{K_{y_c}}{\tau_c^2 s^2 + 2\tau_{y_c}\zeta_{y_c} s + 1} \tag{3.6}$$

where y will be either CPD or $Craving$. These equations suggest a single counseling session has a relatively modest effect on the representative patient (Tobacco Use and Dependence Guideline Panel, 2008; McCarthy *et al.*, 2008b).

Table 3.2: Parameter Values of the Dose-, *Quit*-, and *Stress*-Response Open-Loop Models in Continuous-Time Form for the Representative Patient.

Model	Parameters	Parameter Values
P_{cpd_Q}	$K_Q, \tau_{a1_Q}, \tau_{a2_Q}, \tau_Q, \zeta_Q$	-0.24, 90.53, 10.76, 5.07, 0.59
P_{crav_Q}	$K_Q, \tau_{a1_Q}, \tau_Q, \zeta_Q$	-0.24, 90.53, 5.07, 0.59
P_{cpd_c}	K_c, τ_c, ζ_c	-30.00, 4.00, 1.50
P_{crav_c}	K_c, τ_c, ζ_c	-50.00, 3.75, 1.50
P_{cpd_b}	K_b, τ_b, n_b	-1.28, 0.45, 3.00
P_{crav_b}	K_b, τ_b, n_b	-1.16, 0.50, 3.00
P_{cpd_l}	K_l, τ_{a_l}, τ_l	-0.50, -0.44, 0.88
P_{crav_l}	K_l, τ_{a_l}, τ_l	-0.70, -0.44, 0.50
P_{cpd_s}	K_s, τ_{a_s}, τ_s	1.65, 0.50, 0.80
P_{cpd_s}	K_s, τ_{a_s}, τ_s	3.00, 0.60, 0.80

Bupropion is considered to have a relatively stable and uniform effect when consistently dosed, and has consistently been shown to be an effective cessation aid. A unit increase in daily bupropion dose—here, one 150 mg dose—is thought to take effect within three days (McCarthy *et al.*, 2008b; Piper *et al.*, 2009; Lexicomp, 2014; U.S. Department of Health and Human Services and Centers for Disease Control and Pre-

vention, 2014; Tobacco Use and Dependence Guideline Panel, 2008). Also informed by observations of clinical trial data (McCarthy *et al.*, 2008b), the continuous-time open-loop bupropion dose for the representative subject employs the following structure and the corresponding parameters in Table 3.2:

$$P_{y_b}(s) = \frac{K_{y_b}}{(\tau_b s + 1)^{n_b}} \tag{3.7}$$

Nicotine replacement gum and lozenges deliver nicotine through the bloodstream via tissues in the mouth, which is a significantly slower delivery method compared to actual cigarette smoking. This contributes to the observation reported by patients that a nicotine replacement lozenge dose is less satisfying than one cigarette. However, compared to the daily time scale being considered here, a single dose of lozenge lasts less than one hour generally (Lexicomp, 2014; Piper *et al.*, 2009). Consequently, the gain of the dose-response models to u_l are relatively modest. Furthermore, the continuous-time transfer functions are semi-proper, indicating that lozenges taken on a given day are primarily effective on that same day. Specifically, the models describing how one nicotine replacement lozenge affects CPD and $Craving$ levels employ the continuous-time transfer function structure in equation 3.8 and the corresponding parameters in Table 3.2:

$$P_{y_l}(s) = \frac{K_{y_l}(\tau_{a_{y_l}} s + 1)}{(\tau_{y_l} s + 1)} \tag{3.8}$$

Increased stress or bad mood levels for a given day are associated with relapse on that same day, but not significantly on the following days (Shiffman and Waters, 2004). To reflect this, the disturbance models associated with the exogenous $Stress$ disturbance employ the semi-proper transfer function structure in equation 3.9 and relatively fast speeds of response, documented in Table 3.2:

$$P_{y_S}(s) = \frac{K_{y_S}(\tau_{a_{y_S}} s + 1)}{(\tau_{y_S} s + 1)} \tag{3.9}$$

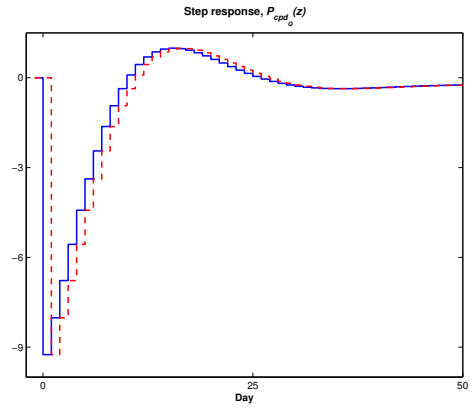

(a) *CPD* Response to *Quit* Unit Step.

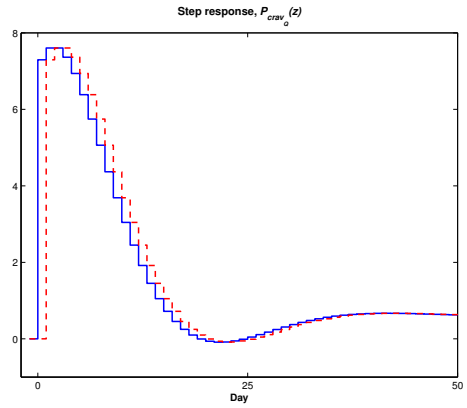

(b) *Craving* Response to *Quit* Unit Step.

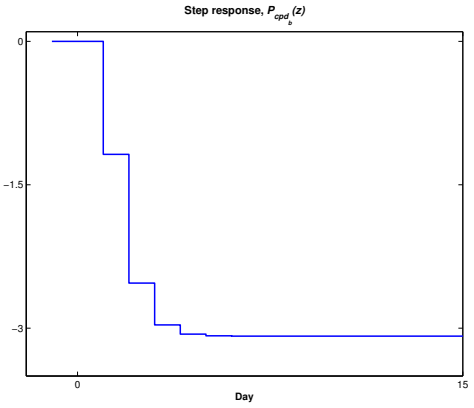

(c) *CPD* Response to u_b Unit Step.

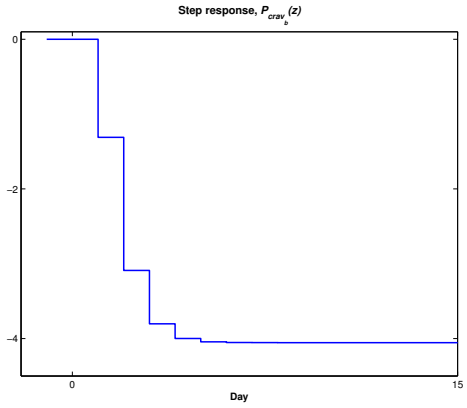

(d) *Craving* Response to u_b Unit Step.

Figure 3.3: Step Responses for the Discrete-Time, Open-Loop Models Representing the Dynamics for a Hypothetical Patient, in Deviation Variable Form (Solid) and the Corresponding Nominal Model if Different than the Hypothetical Patient Model (Dashed). The Unit Step Occurs at $t = 0$.

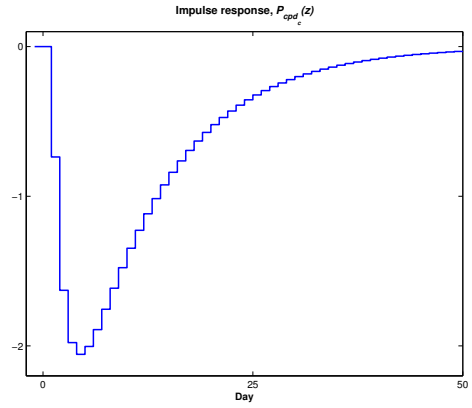

(a) *CPD* Response to u_c Unit Impulse.

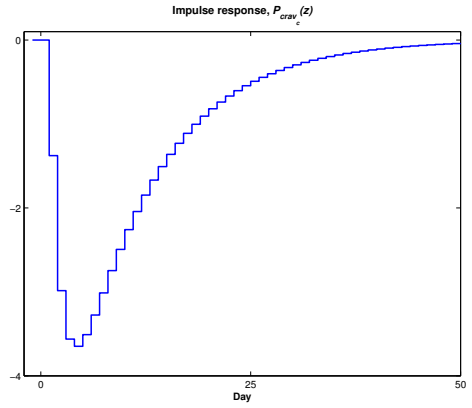

(b) *Craving* Response to u_c Unit Impulse.

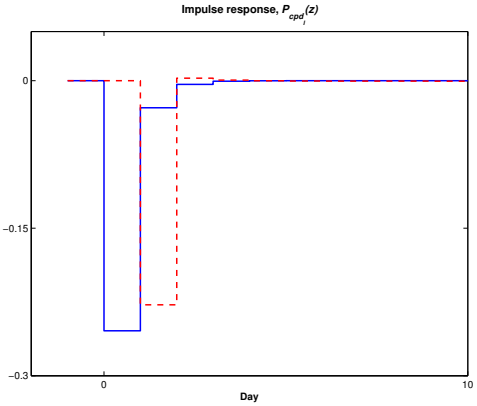

(c) *CPD* Response to u_l Unit Impulse.

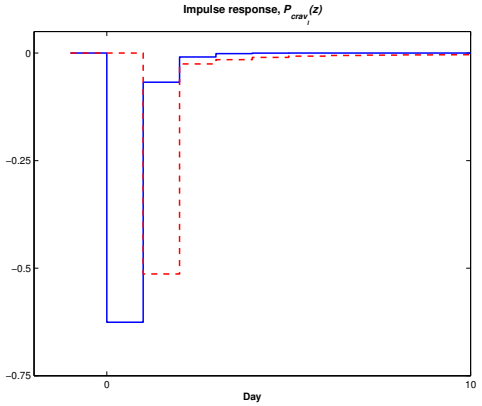

(d) *Craving* Response to u_l Unit Impulse.

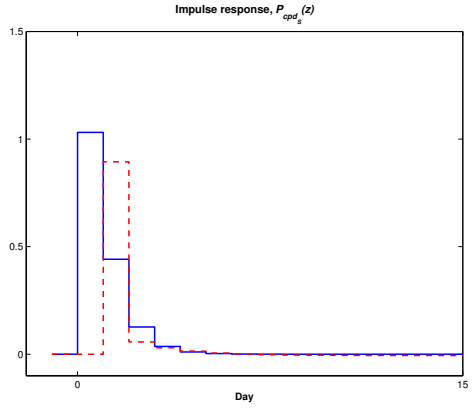

(e) *CPD* Response to *Stress* Unit Impulse.

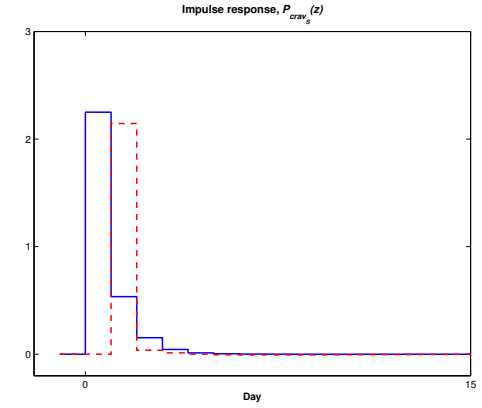

(f) *Craving* Response to *Stress* Unit Impulse.

Figure 3.4: Impulse Responses for the Discrete-Time, Open-Loop Models Representing the Dynamics for a Hypothetical Patient, in Deviation Variable Form (Solid) and the Corresponding Nominal Model if Different than the Hypothetical Patient Model (Dashed). The Unit Impulse Occurs at $t = 0$.

3.3.2 Nominal Models

Section 3.3.1 represents the hypothetical, representative patient's behavior change processes in continuous-time model form. However, the nominal models used in computation of the control action draw from discrete-time representations of these open-loop models. The nominal models are obtained through discretization of equations 3.5 through 3.9 for a 1 day sampling time using the zero-order hold (ZOH) in equation 3.10 (Franklin *et al.*, 1998; The MathWorks, 2014a):

$$G(z) = (1 - z^{-1})\mathcal{Z}\left\{\frac{G(s)}{s}\right\} \tag{3.10}$$

where $G(z)$ is the ZOH-equivalent transfer function representation of $G(s)$, the frequency-domain, continuous-time transfer function being discretized; $\mathcal{Z}\{\}$ indicates the transformation from the continuous-time domain to the discrete-time domain; and z^{-1} indicates a lagged sample (e.g., $z^{-1}y(z)$ in the z-domain is equivalent to $q^{-1}y(k)$ which equals $y(k-1)$ in the time-domain, where q^{-1} is the backshift operator; Franklin *et al.*, 1998; Ljung, 1999). Step or impulse responses for the z-domain versions of the representative subject's open-loop models are depicted with the solid lines in Fig. 3.3 and Fig. 3.4 (step responses are shown for the models in which a step is conceptually the appropriate input, and impulse responses are shown for the models in which an impulse is conceptually the appropriate input). Simulations of the intervention employ state-space versions of the semi-proper, discretized models as the plant. The nominal models describing the response of CPD and $Craving$ to u_c and u_b are discretized versions of equations 3.6 and 3.7 in state-space form.

Discretization of equations 3.5, 3.8, and 3.9 similarly results in semi-proper functions. Correspondingly, the discrete-time state-space form of these equations will feature nonzero direct feedthrough terms. However, open-loop models featuring direct feedthrough terms cannot be integrated into existing HMPC routines in this

form: direct feedthrough terms mean the prediction step would require $u(k+1)$ to compute $y(k+1)$ and so on. Intuitively, this leads to an acausal scenario in the classical HMPC formulation as $u(k+1)$ is itself determined using the predictions $y(k+1)$. Consequently, the controller formulation relies on adjusted versions of discretized equations 3.5, 3.8, and 3.9.

As *Quit* represents a transition from not attempting to quit smoking to attempting to quit, and TQD is typically determined weeks prior to initiation of the attempt (Tobacco Use and Dependence Guideline Panel, 2008; Lexicomp, 2014), this disturbance signal can be constructed *a priori*. Consequently, the *Quit* disturbance signal observed by the controller is defined in equation 3.11:

$$\begin{aligned} Quit &= 0,\ t < \text{TQD} - 1 \\ &= 1,\ t \geq \text{TQD} - 1 \end{aligned} \tag{3.11}$$

The *Quit*-response nominal models can now be appropriately represented by delayed versions of discretized equation 3.5. In other words, as:

$$\tilde{P}_{cpd_Q}(z) = z^{-1} P_{cpd_Q}(z) \tag{3.12}$$

$$\tilde{P}_{crav_Q}(z) = z^{-1} P_{crav_Q}(z) \tag{3.13}$$

where $P_{cpd_Q}(z)$ and $P_{crav_Q}(z)$ are obtained via the transformation described by equation 3.10.

The nominal models for the controlled variables' responses to u_l and $Stress$ are

represented in discrete-time transfer function form as:

$$\tilde{P}_{cpd_l}(z) = \frac{-0.228z^{-1} + 0.227z^{-2}}{1 - 1.359z^{-1} + 0.361z^{-2}} + z^{-1}\frac{0.085z^{-1} - 0.085z^{-2}}{1 - 1.359z^{-1} + 0.361z^{-2}} \tag{3.14}$$

$$\tilde{P}_{crav_l}(z) = \frac{-0.514z^{-1} + 0.510z^{-2}}{1 - 1.542z^{-1} + 0.542z^{-2}} + z^{-1}\frac{0.257z^{-1} - 0.255z^{-2}}{1 - 1.542z^{-1} + 0.542z^{-2}} \tag{3.15}$$

$$\tilde{P}_{cpd_S}(z) = \frac{0.894z^{-1} - 0.990z^{-2}}{1 - 1.570z^{-1} + 0.570z^{-2}} - z^{-1}\frac{0.447z^{-1} - 0.450z^{-2}}{1 - 1.570z^{-1} + 0.570z^{-2}} \tag{3.16}$$

$$\tilde{P}_{crav_S}(z) = \frac{2.145z^{-1} - 1.773z^{-2}}{1 - 1.344z^{-1} + 0.431z^{-2}} - z^{-1}\frac{1.072z^{-1} - 0.887z^{-2}}{1 - 1.344z^{-1} + 0.431z^{-2}} \tag{3.17}$$

These models result from delayed versions of the continuous-time models transformed into the z-domain, but have been adjusted such that the peak effect of a unit impulse in u_l or $Stress$ is not fully "expected" a full day after the impulse occurs, according to the nominal model; furthermore, these models feature minimal mismatch in the effects of a unit impulse in u_l and $Stress$ for two or more days after the unit change. The step and impulse responses for the nominal models are found in Fig. 3.3 and Fig. 3.4.

3.3.3 Dosing Capacity

As described in Section 3.2, the decision framework should feature functionality that systematically balances intervention targets with concerns of unnecessary or over-dosing. Clinically, explicit consideration of total dosing demands may be motivated by concerns for potential side-effects, resource management, or whether a patient is likely to adhere to aggressive dosing schedules.

In production-inventory control literature, a construct representing the quantity of a manipulated variable assigned over a given time frame is referred to as a "capacity" or "work-in-progress" (WIP; Nandola and Rivera, 2013; Schwartz and Rivera, 2010). The capacity for some manipulated variable u is generally represented as:

$$WIP(k+1) = \sum_{i=0}^{\theta-1} u(k-i) \tag{3.18}$$

where θ is the time delay between $u(k)$ and $y(k)$ and $WIP(k+1)$ is the total amount of u assigned by the controller between times k and $k-\theta+1$ (Nandola and Rivera, 2013).

In this context, the capacity construct representing the total amount treatment doses assigned over a given time period is represented by equation 3.19.

$$WIP_T(k+1) = \sum_{j=1}^{n_u} \sum_{i=0}^{n_{wip}} u_j(k-i) \tag{3.19}$$

where j indicates the j^{th} manipulated variable, n_u is the total number of manipulated variables, and n_{wip} is the number of samples in the time frame of interest for the capacity construct.

A metric of the aggregate treatment used each day is defined. This metric of aggregate dosing is calculated according to equation 3.19 where $n_u = 3$, $n_{wip} = 1$, and the u_j's correspond to u_c, u_b, and u_l. It can be shown that this treatment dosing metric is equivalently represented by the discrete-time transfer function,

$$WIP_T(z^{-1}) = \left(\frac{z^{-1}}{1-z^{-1}}\right)\left(u_c(z^{-1}) + u_b(z^{-1}) + u_l(z^{-1})\right) \tag{3.20}$$

which is the summation of u_c, u_b, and u_l dosages assigned for the previous day. For the primary aim of penalizing over-dosing, equation 3.20, or more generally equation 3.19, can be treated as a controlled variable with a set point, objective function penalty weight, and/or constraints (as detailed later in this chapter). WIP_T is assigned a set point equal to 0 in lieu of penalties on each individual treatment component. Such a set point and penalty weight ensures that total dosing demands would be explicitly accounted for during minimization of the objective function while allowing for flexibility in how the treatment components are assigned.

3.4 Controller Development

3.4.1 Overview

Design of the smoking cessation intervention in this chapter draws from an improved formulation of MPC for linear hybrid systems. The HMPC algorithm formulated here considers the behavior change system in a *mixed logical and dynamical* (MLD) framework to systematically manage the general goals of tracking intervention targets, managing measured disturbances, and being robust to unmeasured influences including noise and plant-model mismatch. Ultimately, 3DoF tuning capabilities are incorporated into the decision system to allow more clinician-friendly adjustment to the character of the dosing and outcomes. Details of the formulation are presented below, and the intervention is evaluated through simulation in Chapter 4.

3.4.2 MLD Systems

Although advanced control approaches have largely emerged in the context of systems with smooth dynamics described by smooth transition functions, many systems in which controllers can be incorporated have components that can be described by logic, such as switched systems (Bemporad and Morari, 1999). The discrete-valued nature of the manipulated variable levels in the cessation intervention means the open-loop dynamical system can be represented as an MLD system. MLD systems can be described by linear dynamic equations subject to linear mixed-integer inequalities; these expressions are a function of both continuous and binary variables, with terms for real or integer states, inputs, and constraints. Specifically, the linear hybrid behavior change process is represented as an MLD system in discrete-time, state-space

form as:

$$x(k) = Ax(k-1) + B_1u(k-1) + B_2\delta(k-1) + B_3z(k-1) + B_dd(k-1) \quad (3.21)$$

$$y(k) = Cx(k) + d'(k) + \nu(k) \quad (3.22)$$

$$E_5 \geq E_2\delta(k-1) + E_3z(k-1) - E_4y(k-1) - E_1u(k-1) + E_dd(k-1) \quad (3.23)$$

where x and u represent discrete and continuous states and discrete and continuous inputs, respectively; y is a vector of the outputs; d, d', and ν are the measured disturbances, unmeasured disturbances, and noise, respectively, which are lumped into the measurement equation (equation 3.22); and δ ($\in \{0, 1\}$) and z are discrete and continuous auxiliary variables, which facilitate conversion of logical/discrete decisions into linear inequality constraints (equation 3.23; Bemporad and Morari, 1999).

To illustrate the conversion of logical/discrete decisions into linear inequality constraints via the auxiliary variables, recall equations 3.1 through 3.3. These expressions noted that u_c can only assume one of two possible levels, u_b can only assume one of three possible levels, and u_l can only assume one of 21 possible levels. The discrete-valued nature of counseling dosages can be represented logically:

$$\delta_0(k) = 1 \Leftrightarrow z_0(k) = 0$$

$$\delta_1(k) = 1 \Leftrightarrow z_1(k) = 1 \quad (3.24)$$

$$\sum_{i=0}^{1} \delta_i(k) = 1 \quad (3.25)$$

$$u_c(k) = \sum_{i=0}^{1} z_i(k) \quad (3.26)$$

Equations 3.24 and 3.26 denote that δ_0 represents whether the 0 counseling dose level is "on" and δ_1 represents whether the 1 counseling dose level is "on". Equation 3.25 enforces that $u_c(k)$ must assume one and only one of its possible dosage levels at a time. Similarly, the discrete-valued nature of the bupropion and lozenge dosages can

be represented as equations 3.27 through 3.29 and 3.30 through 3.32, respectively.

$$\delta_j(k) = 1 \Leftrightarrow z_j(k) = j - 2; \quad j \in \{2, 3, 4\} \tag{3.27}$$

$$\sum_{j=2}^{4} \delta_j(k) = 1 \tag{3.28}$$

$$u_b(k) = \sum_{j=2}^{4} z_j(k) \tag{3.29}$$

$$\delta_k(k) = 1 \Leftrightarrow z_k(k) = k - 5; \quad k \in \{5, ..., 25\} \tag{3.30}$$

$$\sum_{k=5}^{25} \delta_k(k) = 1 \tag{3.31}$$

$$u_l(k) = \sum_{k=5}^{25} z_k(k) \tag{3.32}$$

These logical represents are incorporated into the decision-making process via the E_* matrices (where $*$ is 1, 2, 3, 4, 5, and d) in equation 3.23.

Modeled as an exogenous effect in the measurement equation (equation 3.22), the unmeasured disturbance signal d' is considered here to be a stochastic signal that is described by equations 3.33 and 3.34,

$$x_w(k) = A_w x_w(k-1) + B_w w(k-1) \tag{3.33}$$

$$d'(k) = C_w x_w(k) \tag{3.34}$$

where $w(k-1)$ is a vector of integrated white noise; assuming d' consists of uncorrelated components, $B_w = C_w = I$ and $A_w = 0$ (for single-integrating disturbances, i.e., Type I disturbances, as are considered here; Nandola and Rivera, 2013).

The prediction step of HMPC-based decision-making draws from an augmented form of the state-space models in equations 3.21 and 3.22 to represent the system in

difference form:

$$X(k) = \mathcal{A}X(k-1) + \mathcal{B}_1\Delta u(k-1) + \mathcal{B}_2\Delta\delta(k-1) + \mathcal{B}_3\Delta z(k-1) + \mathcal{B}_d\Delta d(k-1) + \mathcal{B}_w\Delta w(k-1) \tag{3.35}$$

$$y(k) = \mathcal{C}X(k) + \nu(k) \tag{3.36}$$

where

$$X(k) = \begin{bmatrix} \Delta x^T(k) \\ \Delta x_w^T(k) \\ y^T(k) \end{bmatrix} = \begin{bmatrix} (x(k) - x(k-1))^T \\ (x_w(k) - x_w(k-1))^T \\ y^T(k) \end{bmatrix} \tag{3.37}$$

$$\mathcal{A} = \begin{bmatrix} A & 0 & 0 \\ 0 & A_w & 0 \\ CA & A_w & I \end{bmatrix}$$

$$\mathcal{B}_i = \begin{bmatrix} B_i \\ 0 \\ CB_i \end{bmatrix},\ i = 1, 2, 3, d\ ;\ \mathcal{B}_w = \begin{bmatrix} 0 \\ I \\ I \end{bmatrix}$$

$$\mathcal{C} = [0 \ \ 0 \ \ I] \tag{3.38}$$

$\Delta*$ denotes $*(k) - *(k-1)$ and $\Delta w(k)$ is a white noise sequence (Nandola and Rivera, 2013).

3.4.3 Prediction Step

The open-loop models in equations 3.35 and 3.36 and the constraints in equation 3.23 are used to predict the outcomes p steps into the future, $\mathcal{Y}(k+1)$,

$$\mathcal{Y}(k+1) = \begin{bmatrix} y^T(k+1) & y^T(k+2) & \cdots & y^T(k+p) \end{bmatrix}^T \tag{3.39}$$

according to

$$\mathcal{Y}(k+1) = \Phi X(k) + \mathcal{H}_1\mathcal{U}(k) + \mathcal{H}_2\bar{\delta}(k) + \mathcal{H}_3\mathcal{Z}(k) + \mathcal{H}_d\mathcal{D}(k) \\ - H_{11}u(k-1) - H_{21}\delta(k-1) - H_{31}z(k-1) - H_{d1}d_{flt}(k-1) \tag{3.40}$$

where the decision variables are represented in equations 3.41 through 3.43, and $\mathcal{D}(k)$ is the externally-forecasted measured disturbance vector, per equation 3.44:

$$\mathcal{U}(k) = \begin{bmatrix} u^T(k) & u^T(k+1) & \cdots & u^T(k+m-1) \end{bmatrix}^T \tag{3.41}$$

$$\bar{\delta}(k) = \begin{bmatrix} \delta^T(k) & \delta^T(k+1) & \cdots & \delta^T(k+p-1) \end{bmatrix}^T \tag{3.42}$$

$$\mathcal{Z}(k) = \begin{bmatrix} z^T(k) & z^T(k+1) & \cdots & z^T(k+p-1) \end{bmatrix}^T \tag{3.43}$$

$$\mathcal{D}(k) = \begin{bmatrix} d_{flt}^T(k) & d_{flt}^T(k+1) & \cdots & d_{flt}^T(k+p-1) \end{bmatrix}^T \tag{3.44}$$

Deriving from equations 3.21 and 3.22, the coefficient matrices Φ, H_{11}, H_{21}, H_{31}, H_{d1}, $\mathcal{H}_1$, $\mathcal{H}_2$, $\mathcal{H}_3$, and $\mathcal{H}_d$ are as follows:

$$\Phi = \begin{bmatrix} \mathcal{C}\mathcal{A} \\ \mathcal{C}\mathcal{A}^2 \\ \vdots \\ \mathcal{C}\mathcal{A}^p \end{bmatrix} \; ; \; H_{i1} = \begin{bmatrix} \mathcal{C}\mathcal{B}_i \\ \mathcal{C}\mathcal{A}\mathcal{B}_i \\ \mathcal{C}\mathcal{A}^2\mathcal{B}_i \\ \vdots \\ \mathcal{C}\mathcal{A}^{p-1}\mathcal{B}_i \end{bmatrix}, \; i = 1, 2, 3, d$$

$$\mathcal{H}_1 = \begin{bmatrix} \mathcal{C}\mathcal{B}_1 & 0 & \cdots & 0 & 0 \\ \mathcal{C}\mathcal{A}\mathcal{B}_1 - \mathcal{C}\mathcal{B}_1 & \mathcal{C}\mathcal{B}_1 & \cdots & 0 & 0 \\ \mathcal{C}\mathcal{A}^2\mathcal{B}_1 - \mathcal{C}\mathcal{A}\mathcal{B}_1 & \mathcal{C}\mathcal{A}\mathcal{B}_1 - \mathcal{C}\mathcal{B}_1 & \ddots & \vdots & \vdots \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathcal{C}\mathcal{A}^{m-1}\mathcal{B}_1 - \mathcal{C}\mathcal{A}^{m-2}\mathcal{B}_1 & \mathcal{C}\mathcal{A}^{m-2}\mathcal{B}_1 - \mathcal{C}\mathcal{A}^{m-3}\mathcal{B}_1 & \cdots & \mathcal{C}\mathcal{A}\mathcal{B}_1 - \mathcal{C}\mathcal{B}_1 & \mathcal{C}\mathcal{B}_1 \\ \mathcal{C}\mathcal{A}^{m}\mathcal{B}_1 - \mathcal{C}\mathcal{A}^{m-1}\mathcal{B}_1 & \mathcal{C}\mathcal{A}^{m-1}\mathcal{B}_1 - \mathcal{C}\mathcal{A}^{m-2}\mathcal{B}_1 & \cdots & \mathcal{C}\mathcal{A}^2\mathcal{B}_1 - \mathcal{C}\mathcal{A}\mathcal{B}_1 & \mathcal{C}\mathcal{A}\mathcal{B}_1 \\ \mathcal{C}\mathcal{A}^{m}\mathcal{B}_1 - \mathcal{C}\mathcal{A}^{m-1}\mathcal{B}_1 & \mathcal{C}\mathcal{A}^{m-1}\mathcal{B}_1 - \mathcal{C}\mathcal{A}^{m-2}\mathcal{B}_1 & \cdots & \mathcal{C}\mathcal{A}^2\mathcal{B}_1 - \mathcal{C}\mathcal{A}\mathcal{B}_1 & \mathcal{C}\mathcal{A}\mathcal{B}_1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathcal{C}\mathcal{A}^{p-1}\mathcal{B}_1 - \mathcal{C}\mathcal{A}^{p-2}\mathcal{B}_1 & \mathcal{C}\mathcal{A}^{p-2}\mathcal{B}_1 - \mathcal{C}\mathcal{A}^{p-1}\mathcal{B}_1 & \cdots & \mathcal{C}\mathcal{A}^{p-m+1}\mathcal{B}_1 - \mathcal{C}\mathcal{A}^{p-m}\mathcal{B}_1 & \mathcal{C}\mathcal{A}^{p-m}\mathcal{B}_1 \end{bmatrix}$$

$$\mathcal{H}_i = \begin{bmatrix} \mathcal{C}\mathcal{B}_i & 0 & \cdots & 0 & 0 \\ \mathcal{C}\mathcal{A}\mathcal{B}_i - \mathcal{C}\mathcal{B}_i & \mathcal{C}\mathcal{B}_i & \cdots & 0 & 0 \\ \mathcal{C}\mathcal{A}^2\mathcal{B}_i - \mathcal{C}\mathcal{A}\mathcal{B}_i & \mathcal{C}\mathcal{A}\mathcal{B}_i - \mathcal{C}\mathcal{B}_i & \cdots & \vdots & \vdots \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathcal{C}\mathcal{A}^{p-1}\mathcal{B}_i - \mathcal{C}\mathcal{A}^{p-2}\mathcal{B}_i & \mathcal{C}\mathcal{A}^{p-2} - \mathcal{C}\mathcal{A}^{p-3}\mathcal{B}_i & \cdots & \mathcal{C}\mathcal{A}\mathcal{B}_i - \mathcal{C}\mathcal{B}_i & \mathcal{C}\mathcal{B}_i \end{bmatrix}, \quad i = 2, 3, d$$

The inequality expression over the prediction horizon propagates to,

$$\begin{aligned}\mathcal{E}_5 \geq \mathcal{E}_2\bar{\delta}(k) + \mathcal{E}_3\mathcal{Z}(k) + \mathcal{E}_1\mathcal{U}(k) + \mathcal{E}_4 X(k) + \mathcal{E}_d\mathcal{D}(k) \\ - \mathcal{E}_{41}u(k-1) - \mathcal{E}_{42}\delta(k-1) - \mathcal{E}_{43}z(k-1) - \mathcal{E}_{4d}d_{flt}(k-1) \end{aligned} \tag{3.45}$$

where

$$\mathcal{E}_i = (\bar{E}_4\bar{\mathcal{H}}_i + \bar{E}_i),\ i = 2, 3, d$$

$$\mathcal{E}_4 = \bar{E}_4\bar{\Phi}$$

$$\mathcal{E}_{4i} = \bar{E}_4\bar{\mathcal{H}}_{i1},\ i = 1, 2, 3, d$$

$$\mathcal{E}_5 = [E_5 \ \ E_5 \ \ \cdots \ \ E_5]^T$$

and

$$\bar{E}_1 = \begin{bmatrix} -E_1 & 0 & \cdot & 0 \\ 0 & \ddots & \cdot & \vdots \\ \vdots & \cdots & \cdot & -E_1 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & \cdot & -E_1 \end{bmatrix}$$

$$\bar{E}_i = \text{diag}\{E_i, \cdots, E_i\},\ i = 2, 3, d$$

$$\bar{E}_4 = \text{diag}\{-E_4, \cdots, -E_4\}$$

$$\bar{\mathcal{H}}_j = \begin{bmatrix} [0]_{n_y} \\ \mathcal{H}_j(1:(p-1)n_y, :) \end{bmatrix},\ j = 1, 2, 3, d, 11, 21, 31, d1$$

$$\bar{\Phi} = \begin{bmatrix} \mathcal{C} \\ \Phi(1:(p-1)n_y, :) \end{bmatrix}$$

3.4.4 Objective Function, Constraints, & Targets

The optimization problem at each decision period consists of determining the sequence of control actions $\{u(k), \cdots, u(k+m-1)\}$, $\{\delta(k), \cdots, \delta(k+p-1)\}$, and

$\{z(k), \cdots, z(k+p-1)\}$ that minimize the value of the objective function J,

$$\min_{\{[u(k+i)]_{i=0}^{m-1},[\delta(k+i)]_{i=0}^{p-1},[z(k+i)]_{i=0}^{p-1}\}} J \tag{3.46}$$

subject to process constraints, where:

$$J \triangleq \sum_{i=1}^{p} ||y(k+i) - y_r(k+i)||^2_{Q_y} + \sum_{i=0}^{m-1} ||\Delta u(k+i)||^2_{Q_{\Delta u}} + \sum_{i=0}^{m} ||u(k+i) - u_r(k+i)||^2_{Q_u} + \sum_{i=0}^{p-1} ||\delta(k+i) - \delta_r||^2_{Q_\delta} + \sum_{i=0}^{p-1} ||z(k+i) - z_r||^2_{Q_z} \tag{3.47}$$

$(\cdot)_r$ denotes a reference signal and $||(\cdot)||_{Q_*} \triangleq \sqrt{(\cdot)^T \mathbf{Q}_* (\cdot)}$ is the vector 2-norm weighted by the penalty matrix $\mathbf{Q}_*$: Q_y is control error penalty weight; $Q_{\Delta u}$ is the move size penalty weight; Q_u penalizes deviations of $u(k+i)$ from $u_r(k+i)$; Q_δ penalizes deviations of $\delta(k+i)$ from δ_r; and Q_z penalizes deviations of $z(k+i)$ from z_r. The process constraints corresponding to the quadratic cost function in equation 3.47 are captured in equations 3.23 and,

$$y_{min} \leq y(k+i) \leq y_{max}\,,\ 1 \leq i \leq p \tag{3.48}$$

$$u_{min} \leq u(k+i) \leq u_{max}\,,\ 0 \leq i \leq m-1 \tag{3.49}$$

$$\Delta u_{min} \leq \Delta u(k+i) \leq \Delta u_{max}\,,\ 0 \leq i \leq m-1 \tag{3.50}$$

where y_{min} and y_{max} are the upper and lower bounds on the controlled variable trajectory, u_{min} and u_{max} are the bounds on the manipulated variables, and Δu_{min} and Δu_{max} are the bounds on the manipulated variable move sizes.

Intervention Objective Function

Corresponding to Fig. 3.1 and equation 3.47, the decision system developed here focuses on the general problem of minimizing equation 3.51,

$$\begin{aligned} J \triangleq\ & ||CPD(k+i) - CPD_r(k+i)||^2_{Q_{cpd}} + ||Craving(k+i) - Craving_r(k+i)||^2_{Q_{crav}} \\ & + ||\Delta u_c(k+i)||^2_{Q_{\Delta u_c}} + ||\Delta u_b(k+i)||^2_{Q_{\Delta u_b}} + ||\Delta u_c(k+i)||^2_{Q_{\Delta u_l}} \\ & + ||u_c(k+i)||^2_{Q_{u_c}} + ||u_b(k+i)||^2_{Q_{u_b}} + ||u_l(k+i)||^2_{Q_{u_l}} + \cdots \end{aligned} \tag{3.51}$$

where Q_{cpd} is the objective function penalty weight for the CPD tracking error and Q_{crav} is that for the $Craving$ tracking error; $Q_{\Delta u_c}$, $Q_{\Delta u_b}$, and $Q_{\Delta u_l}$ are move suppression weights for the three treatment components; Q_{u_c}, Q_{u_b}, and Q_{u_l} are the penalty weights for use of the individual treatment components. Note, equation 3.51 assumes the treatment component set points are 0 for all times, and omits terms for the auxiliary variables and capacities for brevity.

Intervention Constraints

Directly related to the constraints described in equations 3.1 through 3.32 are the following minimum and maximum manipulated variable bounds:

$$0 \leq u_c(k) \leq 1 \tag{3.52}$$

$$0 \leq u_b(k) \leq 2 \tag{3.53}$$

$$0 \leq u_l(k) \leq 20 \tag{3.54}$$

In terms of move size constraints, there is no limitation on the degree to which counseling and lozenge dose can be adjusted from one day to another, beyond those

implied by equations 3.52 and 3.54:

$$-1 \leq \Delta u_c(k) \leq 1 \tag{3.55}$$

$$-20 \leq \Delta u_l(k) \leq 20 \tag{3.56}$$

A patient cannot jump directly from zero to two 150 mg doses of bupropion per day, instead, they must increase by one dose at a time. However, a patient at the maximum bupropion dose can reduce their dose to any level (Lexicomp, 2014). These dosing factors translate directly to the following move size constraint:

$$-2 \leq \Delta u_b(k) \leq 1 \tag{3.57}$$

However, $u_b = 2$ can be assigned only after $u_b = 1$ for at least three days (Lexicomp, 2014). Capturing the fact that going from one to two bupropion doses can occur after a minimum of three days but at any time after that would require the decision-making computations account for a "days-since-last-u_b-change" component. Incorporating a time-dependent constraint of this nature involves a significantly more complex task than the computational intensity associated with constraints that are non-time-varying. Instead, the three day ramp up period is incorporated as a switching time restriction. Specifically,

$$\begin{aligned} 0 \leq \Delta u_b(k) \leq 0,\ k \neq\ t_{sw}^1 + nT_{sw} \\ -2 \leq \Delta u_b(k) \leq 1,\ k =\ t_{sw}^1 + nT_{sw},\ n = 0, 1, 2, ... \end{aligned} \tag{3.58}$$

where t_{sw}^1 is the first decision period at which u_b dose can be adjusted. T_{sw} is the switching time. $i = 2$ and $T_{sw} = 4$ are primarily considered in the following and correspond to possible u_b adjustments in the week pre-TQD that closely reflects current clinical practice in terms of bupropion dosing (McCarthy *et al.*, 2008b). Additional details on switching time considerations within the context of HMPC-based decision frameworks for health behavior therapeutics can be found in Dong *et al.* (2014).

The following logical bounds constitute the constraint in equation 3.48 in the day-to-day decision-making considered here:

$$0 \leq CPD(k+1) \leq \infty \tag{3.59}$$

$$0 \leq Craving(k+1) \leq \infty \tag{3.60}$$

To reflect the fact that the total amount of counseling is limited, a capacity construct is employed (Nandola and Rivera, 2013). Drawing from equation 3.18, the controller calculates the total amount of counseling assigned from the beginning of the intervention through the current day as:

$$WIP_c(k+1) = \sum_{i=0}^{k} u_c(k-i) \tag{3.61}$$

With this construct, counseling can be assigned on any day, but it can be limited to only a certain number of sessions over the entire duration of the intervention. For a maximum of five counseling sessions,

$$0 \leq WIP_c(k+1) \leq 5 \tag{3.62}$$

To facilitate dosing that takes advantage of a synergistic effect between bupropion and lozenge, an additional linear inequality is defined to ensure that lozenge doses are only assigned when $u_b(k) \neq 0$; this constraint is represented logically as:

$$\delta_m = 1\,;\ m \in \{3,4,5\} \tag{3.63}$$

$$\sum_{m=3}^{5} \delta_m(k) = 1 \tag{3.64}$$

Intervention Targets

Intuitively, TQD corresponds to the time at which the CPD and $Craving$ targets change; mathematically:

$$\begin{aligned} CPD_{target}(k) &= CPD_{base}\,,\ t < \text{TQD} \\ &= 0\,, \qquad\qquad t \geq \text{TQD} \end{aligned} \tag{3.65}$$

$$\begin{aligned} Craving_{target}(k) &= Craving_{base}\,,\ t < \text{TQD} \\ &= 0\,, \qquad\qquad t \geq \text{TQD} \end{aligned} \tag{3.66}$$

where CPD_{base} and $Craving_{base}$ are the pre-TQD levels of CPD and $Craving$. Equation 3.65 reflects that the intervention essentially seeks to switch smoking from "on" pre-TQD to "off" as of TQD. As TQD is assumed to be defined by a clinician or a patient at least two weeks prior to actual initiation of the quit attempt (Tobacco Use and Dependence Guideline Panel, 2008; McCarthy *et al.*, 2008b; Lexicomp, 2014), the reference vector $\mathcal{Y}_r$ in equation 3.47 is essentially a function of TQD, and is recalculated at each decision time to reflect the set point change in equations 3.65 and 3.66.

3.4.5 Solving the Optimization Problem

Drawing from equations 3.40 and 3.45, the optimization problem described in equations 3.46 and 3.47 can be rewritten in vector form as,

$$\begin{aligned} \min_{\{\mathcal{U}(k),\bar{\delta}(k),\mathcal{Z}(k)\}} J \triangleq\ & ||\mathcal{Y}(k+1) - \mathcal{Y}_r||^2_{\widehat{Q}_y} + ||R_u\mathcal{U}(k) - R_{u0}u(k-1)||^2_{\widehat{Q}_{\Delta u}} \\ & + ||\mathcal{U}(k) - \mathcal{U}_r||^2_{\widehat{Q}_u} + ||\bar{\delta}(k) - \bar{\delta}_r||^2_{\widehat{Q}_\delta} + ||\mathcal{Z}(k) - \mathcal{Z}_r||^2_{\widehat{Q}_z} \end{aligned} \tag{3.67}$$

subject to equation 3.45 and,

$$\mathcal{Y}_{min} \leq \mathcal{Y}(k+1) \leq \mathcal{Y}_{max} \tag{3.68}$$

$$\mathcal{U}_{min} \leq \mathcal{U}(k) \leq \mathcal{U}_{max} \tag{3.69}$$

$$\Delta\mathcal{U}_{min} \leq \Delta\mathcal{U}(k) \leq \Delta\mathcal{U}_{max} \tag{3.70}$$

where

$$\mathcal{Y}_r = \begin{bmatrix} y_r^T(k+1) & y_r^T(k+2) & \cdots & y_r^T(k+p) \end{bmatrix}^T \tag{3.71}$$

$$\mathcal{U}_r = \begin{bmatrix} u_r^T(k) & u_r^T(k+1) & \cdots & u_r^T(k+m-1) \end{bmatrix}^T \tag{3.72}$$

$$\bar{\delta}_r = \begin{bmatrix} \delta_r^T(k) & \delta_r^T(k+1) & \cdots & \delta_r^T(k+p-1) \end{bmatrix}^T \tag{3.73}$$

$$\mathcal{Z}_r = \begin{bmatrix} z_r^T(k) & z_r^T(k+1) & \cdots & z_r^T(k+p-1) \end{bmatrix}^T \tag{3.74}$$

are the reference vectors and

$$R_u = \begin{bmatrix} I & 0 & 0 & \cdots & 0 & 0 \\ -I & I & 0 & \cdots & \vdots & \vdots \\ 0 & -I & I & \ddots & \vdots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & \cdots & -I & I \end{bmatrix}, \quad R_{u0} = \begin{bmatrix} I \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

$\widehat{Q}_y$, $\widehat{Q}_{\Delta u}$, $\widehat{Q}_u$, $\widehat{Q}_\delta$, and $\widehat{Q}_z$ in equation 3.67 correspond to the penalty weights for the control error, move suppression, treatment components, discrete auxiliary variables, and continuous auxiliary variables, respectively; $\widehat{Q}_*$ indicates a square matrix with the respective penalty weights in the diagonal.

The task at each decision period becomes to determine $\mathcal{U}(k)$, $\bar{\delta}(k)$, and $\mathcal{Z}(k)$ by solving the problem in equation 3.67 subject to the constraints in equation 3.45 and 3.68 through 3.70. In this HMPC framework, this optimization problem consists of a mixed integer quadratic program (MIQP). Specifically, the standard MIQP consists

of,

$$\min_{\xi} J \triangleq \frac{1}{2}\xi^T \mathcal{H}\xi + \mathcal{G}^T\xi \tag{3.75}$$

$$\mathcal{S}\xi \leq b \tag{3.76}$$

where ξ is the vector of decision variables, i.e.,

$$\xi = \begin{bmatrix} \mathcal{U}(k)^T \\ \bar{\delta}(k)^T \\ \mathcal{Z}(k)^T \end{bmatrix} \tag{3.77}$$

$\mathcal{H}$ and $\mathcal{G}$ are the coefficient matrices for the quadratic and linear terms in the objective function, respectively. These terms are obtained after substituting the prediction equation $\mathcal{Y}(k+1)$ (equation 3.40) into the objective function in vector form (equation 3.67) and rearranging to group the quadratic and linear terms. Specifically,

$$\mathcal{H} = 2\begin{bmatrix} \mathcal{H}_1^T\widehat{Q}_y\mathcal{H}_1 + \widehat{Q}_{\Delta u} + \widehat{Q}_Q & \mathcal{H}_1^T\widehat{Q}_y\mathcal{H}_2 & \mathcal{H}_1^T\widehat{Q}_y\mathcal{H}_3 \\ \mathcal{H}_2^T\widehat{Q}_y\mathcal{H}_1 & \mathcal{H}_2^T\widehat{Q}_y\mathcal{H}_2 + \widehat{Q}_\delta & \mathcal{H}_2^T\widehat{Q}_y\mathcal{H}_3 \\ \mathcal{H}_3^T\widehat{Q}_y\mathcal{H}_1 & \mathcal{H}_3^T\widehat{Q}_y\mathcal{H}_2 & \mathcal{H}_3^T\widehat{Q}_y\mathcal{H}_3 + \widehat{Q}_z \end{bmatrix} \tag{3.78}$$

$$\mathcal{G} = 2[g_1 \; g_2 \; g_3]^T \tag{3.79}$$

where

$$
\begin{aligned}
g_1 &= X(k)^T\Phi^T\widehat{Q}_y\mathcal{H}_1 - \mathcal{Y}_r^T\widehat{Q}_y\mathcal{H}_1 - \mathcal{U}_r^T\widehat{Q}_u + \mathcal{D}(k)^T\mathcal{H}_d^T\widehat{Q}_y\mathcal{H}_1 \\
&\quad - u(k-1)^T(R_{u0}^T\widehat{Q}_{\Delta u}R_u + H_{11}^T\widehat{Q}_y\mathcal{H}_1) - \delta(k-1)^T H_{21}^T\widehat{Q}_y\mathcal{H}_1 - z(k-1)^T H_{31}^T\widehat{Q}_y\mathcal{H}_1 \\
&\quad - d_{flt}(k-1)^T H_{d1}^T\widehat{Q}_y\mathcal{H}_1 \\
g_2 &= X(k)^T\Phi^T\widehat{Q}_y\mathcal{H}_2 - \mathcal{Y}_r^T\widehat{Q}_y\mathcal{H}_2 - \bar{\delta}_r^T\widehat{Q}_\delta + \mathcal{D}(k)^T\mathcal{H}_d^T\widehat{Q}_y\mathcal{H}_2 \\
&\quad - u(k-1)^T H_{11}^T\widehat{Q}_y\mathcal{H}_2 - \delta(k-1)^T H_{21}^T\widehat{Q}_y\mathcal{H}_2 - z(k-1)^T H_{31}^T\widehat{Q}_y\mathcal{H}_2 \\
&\quad - d_{flt}(k-1)^T H_{d1}^T\widehat{Q}_y\mathcal{H}_2 \\
g_3 &= X(k)^T\Phi^T\widehat{Q}_y\mathcal{H}_3 - \mathcal{Y}_r^T\widehat{Q}_y\mathcal{H}_3 - \mathcal{Z}_r^T\widehat{Q}_z + \mathcal{D}(k)^T\mathcal{H}_d^T\widehat{Q}_y\mathcal{H}_3 \\
&\quad - u(k-1)^T H_{11}^T\widehat{Q}_y\mathcal{H}_3 - \delta(k-1)^T H_{21}^T\widehat{Q}_y\mathcal{H}_3 - z(k-1)^T H_{31}^T\widehat{Q}_y\mathcal{H}_3 \\
&\quad - d_{flt}(k-1)^T H_{d1}^T\widehat{Q}_y\mathcal{H}_3
\end{aligned}
$$

$\mathcal{S}$ and b in equation 3.76 are the coefficient matrices for the linear constraints. Specifically,

$$\mathcal{S} = [s_1 \ s_2 \ -s_2]^T; \tag{3.80}$$

where

$$s_1 = [\mathcal{E}_1 \ \mathcal{E}_2 \ \mathcal{E}_3]; \quad s_2 = \begin{bmatrix} \mathcal{H}_1 & \mathcal{H}_2 & \mathcal{H}_3 \\ I_{m(nu)} & [0]_{m(nu)xp(n\delta)} & [0]_{m(nu)xp(nz)} \\ R_u & [0]_{m(nu)xp(n\delta)} & [0]_{m(nu)xp(nz)} \end{bmatrix} \tag{3.81}$$

and

$$b = \begin{bmatrix} \mathcal{E}_5 - \mathcal{E}_4 X(k) - \mathcal{E}_d \mathcal{D}(k) + \mathcal{E}_{41} u(k-1) + \mathcal{E}_{42} \delta(k-1) + \mathcal{E}_{43} z(k-1) + \mathcal{E}_{4d} d_{flt}(k-1) \\ \mathcal{Y}_{max} - \Phi X(k) + H_d \mathcal{D}(k) - H_{11} u(k-1) - H_{21} \delta(k-1) - H_{31} z(k-1) - H_{d1} d_{flt}(k-1) \\ \mathcal{U}_{max} \\ \Delta \mathcal{U}_{max} + R_{u0} u(k-1) \\ -\mathcal{Y}_{min} \Phi X(k) + H_d \mathcal{D}(k) - H_{11} u(k-1) - H_{21} \delta(k-1) - H_{31} z(k-1) - H_{d1} d_{flt}(k-1) \\ -\mathcal{U}_{min} \\ -\Delta \mathcal{U}_{min} - R_{u0} u(k-1) \end{bmatrix} \tag{3.82}$$

Well-established commercially available programs can solve this MIQP. The simulations presented in Chapter 4 rely on the TOMLAB-CPLEX solver described in Holmstrom *et al.* (2009).

Output Constraint Relaxation

As previously described, the trajectory of the CPD target is meant to reflect the intent to switch from a state of smoking "on" to smoking "off" (see equation 3.65). Because of this, the for CPD target during the quit attempt is equal to zero cigarettes per day, which is also the physically realizable bound of CPD. In other words, $y_{target}(k) = y_{min} = 0$ for $k \geq$ TQD. As the behavior change process is represented as a linear system here, dosing prior to or around TQD can lead to predicted CPD trajectories that violate the y_{min} constraint. Consequently, early simulations of the intervention frequently resulted in infeasible solutions to the optimization problem according to the TOMLAB-CPLEX solver. The proof-of-concept nature of this work motivated incorporation of constraint-relaxing capabilities, ultimately allowing valuable analysis of the intervention through simulation with often minor relaxation of the y_{min} constraint. Three relaxation approaches were explored.

Initially, a relaxation approach was considered where the window over which a

constraint is enforced is shortened in order to obtain feasibility (Scokaert and Clarke, 1994). Here, this entailed altering the constraint in equation 3.48 to,

$$-\infty \leq y(k+i) \leq y_{max},\ 1 \leq i \leq p_{\epsilon} \tag{3.83}$$

$$y_{min} \leq y(k+i) \leq y_{max},\ p_{\epsilon}+1 \leq i \leq p \tag{3.84}$$

as necessary. Equation 3.83 reflects that the constraint corresponding to the lower bound on the controlled variables is removed for the first p_{ϵ} time points in the prediction horizon; the lower bound constraint is enforced for the duration of the prediction horizon, per equation 3.84. (Note, both bounds could be relaxed in this manner. Equations 3.83 and 3.84 only indicate relaxation of the lower bound as $y_{max} = \infty$ in this intervention already, per equations 3.59 and 3.60.) Equations 3.83 and 3.84 would only be implemented when the TOMLAB-CPLEX solver indicated an infeasible solution under the conditions of the hard constraint in equation 3.48; when relaxation was required, the bounds in equations 3.83 and 3.84 were incorporated in an iterative manner such that the smallest p_{ϵ} that gave feasible solutions in the optimization problem was found.

Alternatively, relaxation was considered in which a slack variable, ϵ, was added to b in equation 3.76, where ϵ is a column vector with the same number of rows as in b. An informal implementation of this slack variable approach was briefly implemented in a loop around the decision computations such that $\epsilon = 0$ initially, and was increased by an operator-defined increment (e.g., $\Delta\epsilon = 0.1$) in an iterative manner until an ϵ value that gave a feasible solution to the optimization problem was found.

A more formal implementation of the slack variable concept was determined to be a more appropriate relaxation method. In this more formal approach, the values of the elements in the column vector ϵ are found optimally (Camacho and Bordons, 1995; Zheng and Morari, 1995). Specifically, the optimization problem described by

equations 3.75 and 3.76 is now,

$$\min_{\xi^*} J \triangleq \frac{1}{2}\xi^{*T}\begin{bmatrix}\mathcal{H} & 0\\ 0 & \widehat{Q}_\epsilon\end{bmatrix}\xi^* + \begin{bmatrix}\mathcal{G}\\ 0\end{bmatrix}^T \xi^* \tag{3.85}$$

$$[\mathcal{S} \ -1]\begin{bmatrix}\xi\\ \epsilon\end{bmatrix} \leq b \tag{3.86}$$

where ξ, $\mathcal{H}$, $\mathcal{G}$, $\mathcal{S}$ and b are still described by equations 3.77, 3.78, 3.79, 3.80 and 3.82, respectively; the vector of decision variables, ξ^*, is now,

$$\xi^* = \begin{bmatrix}\xi\\ \epsilon\end{bmatrix} = \begin{bmatrix}\mathcal{U}(k)^T\\ \bar{\delta}(k)^T\\ \mathcal{Z}(k)^T\\ \epsilon\end{bmatrix} \tag{3.87}$$

and $\widehat{Q}_\epsilon$ is a diagonal matrix of weights penalizing the degree of constraint relaxation.

The $\widehat{Q}_\epsilon$ matrix results from inclusion of a constraint relaxation term in the objective function (which is how this relaxation method is more flexible and optimal than some alternative approaches):

$$\begin{aligned}\min_{\{\mathcal{U}(k),\bar{\delta}(k),\mathcal{Z}(k),\epsilon\}} J \triangleq\ & ||\mathcal{Y}(k+1) - \mathcal{Y}_r||^2_{\widehat{Q}_y} + ||R_u\mathcal{U}(k) - R_{u0}u(k-1)||^2_{\widehat{Q}_{\Delta u}} \\ & + ||\mathcal{U}(k) - \mathcal{U}_r||^2_{\widehat{Q}_u} + ||\bar{\delta}(k) - \bar{\delta}_r||^2_{\widehat{Q}_\delta} + ||\mathcal{Z}(k) - \mathcal{Z}_r||^2_{\widehat{Q}_z} + ||\epsilon||^2_{\widehat{Q}_\epsilon}\end{aligned} \tag{3.88}$$

The values in the diagonal of $\widehat{Q}_\epsilon$ influence how easily solutions to the optimization problems can diverge from the specified hard constraint. Simulations of the intervention here defined $\widehat{Q}_\epsilon$ as:

$$\widehat{Q}_\epsilon = Q_\epsilon\, I \tag{3.89}$$

where Q_ϵ is a single value chosen by the user. Mitigation of unnecessary relaxation motivates choice of the largest Q_ϵ that provides feasible solutions for an acceptable computing time, theoretically. Practically, one should choose $Q_\epsilon \gg \max(Q_y, Q_{\Delta u}, Q_u, Q_\delta, Q_z)$. Unless otherwise specified, $Q_\epsilon = 1e5$ in this dissertation.

The constraint-relaxation method described by equations 3.85 through 3.88 reflects that here, the slack variable is calculated during every decision period, but is constant over the prediction horizon considered in those computations. Furthermore, as the lower bound on the controlled variables were the sources of infeasibility errors from the TOMLAB-CPLEX routines, the rows of ϵ corresponding to the y_{min} rows of equation 3.82 were optimally determined, while the other elements of ϵ are set to 0.

3.5 Three-Degree-of-Freedom (3DoF) Tuning Capabilities

The formulation described in Section 3.4 that is to be incorporated into the general structure in Fig. 3.1 relies on p, m, and the control error, move suppression, and treatment component penalty weights to influence the character of the manipulated variable assignments, controlled variable performance, and overall robustness. However, the effect of adjustments to these parameters on individual controlled and manipulated variables is difficult to disentangle.

3DoF tuning capabilities offer a more intuitive way to tune the controller in an attempt to obtain favorable outcome responses and manipulated variable adjustments. More straightforward tuning capabilities could facilitate a more clinician-friendly HMPC-based intervention. Portions of Chapter 4 examine an intervention formulation with 3DoF features. Specifically, the more basic HMPC structure depicted in Fig. 3.1 is expanded to consider the intervention framework depicted in Fig. 3.5.

In this figure, the CPD target, $Craving$ target, and $Stress$ disturbance signals pass through filters that influence tuning. Detailed below, first order filters that

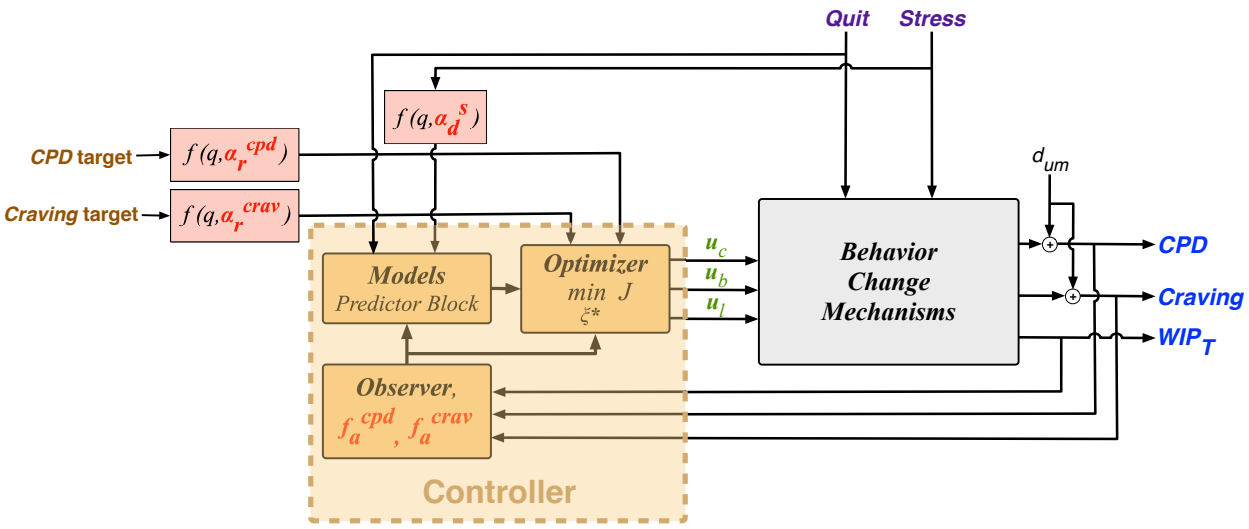

Figure 3.5: Block Diagram of the Decision Framework for an Adaptive, Smoking Cessation Intervention that Employs an HMPC Algorithmic Structure with 3DoF Tuning Functionality.

adhere to the structure in equation 3.91 are considered. α_r^{cpd} and α_r^{crav} influence the speed at which the intervention tracks the CPD and $Craving$ targets ($\alpha_r^* = 0$ corresponds to the most aggressive tracking of the target for a given set of penalty weights and constraints, where $*$ indicates CPD or $Craving$). α_d^S ($\epsilon\ [0, 1)$) affects the character of set point tracking by influencing the speed at which a $Stress$ disturbance is rejected ($\alpha_d^S = 0$ corresponds to the most aggressive $Stress$ disturbance rejection case). The f_a terms influence the speed at which unmeasured disturbances are rejected.

3.5.1 Reference Trajectories

The α_r parameter is intended to facilitate tuning of controlled variable responses independently (Nandola and Rivera, 2013).

In the 3DoF structure, the reference signal(s) incorporated into decision computations (see equation 3.71 and the related equation development) consist of a filtered

representation of the intervention targets, per equations 3.90:

$$\frac{y_r(k+i)}{y_{target}(k+i)} = f(q^{-1}, \alpha_r^j), \;\; 1 \leq j \leq n_j, \;\; 1 \leq i \leq p \tag{3.90}$$

where $f(q^{-1}, \alpha_r^j)$ is a discrete-time filter for reference signal j in the set of n_j targets ($n_j = n_y$); q^{-1} is the backshift operator (Lee and Yu, 1994; Nandola and Rivera, 2013). As the intervention targets in the cessation intervention are steps (equations 3.65 and 3.66), $f(q^{-1}, \alpha_r^j)$ takes the form of a Type-I filter:

$$f(q^{-1}, \alpha_*^j) = \frac{\alpha_* - 1}{\alpha_* q^{-1} - 1}, \;\; 0 \leq \alpha_*^j < 1, \;\; 1 \leq j \leq n_j \tag{3.91}$$

where $*$ denotes r and n_j is still the number of reference signals when the Type-I filter corresponds to equation 3.90. The denominator of the filter in equation 3.91 indicates that filtering the y_{target} signal influences the speed at which the target is tracked, as the signal is processed by a filter that features a pole on the real axis, the location of which is between 0 and -1. The parameter α_r^j individually detunes the speed at which the intervention target corresponding to reference (equivalently, outcome) j is tracked: $\alpha_r^j = 0$ corresponds to the no-detuning case, in which the j^{th} target signal exactly equals the corresponding y_r values in equations 3.71 at all time points. In difference equation form, equation 3.91 in the context of the reference signal filter expands to,

$$y_r(k+i) = (1 - \alpha_r)\, y_{target}(k+i) + \alpha_r y_r(k+i-1) \tag{3.92}$$

Equation 3.92 highlights that α_r generally influences the fraction of the set point signal used in the optimization problem at time $k+i$ that is the actual intervention target at that time.

Note that Fig. 3.5 only depicts filtering action for the CPD and $Craving$ references. Because the WIP_T construct is incorporated in order to discourage aggregate

over-dosing, its target remains constant—equal to 0—over all time points. Consequently, an $\alpha_r^{wip_T}$ tuning knob should play no effective role in dosing decisions, and therefore this target signal is omitted from the block diagram, and no filter is associated with WIP_T. Similarly, the set points for the decision variables are also set to 0 for every time point (equations 3.72 through 3.74 are vectors of 0's), and therefore no filters are associated with these set points.

3.5.2 Measured Disturbances

Equation 3.93 corresponds to the filter that facilitates independent adjustment of the speed at which the j_d^{th} disturbance is rejected.

$$\frac{d_{flt}(k+i)}{d(k+i)} = f(q^{-1}, \alpha_d^{j_d}), \;\; 1 \leq j_d \leq n_d, \;\; 1 \leq i \leq p \tag{3.93}$$

where n_d is the number of measured disturbances and $d(k+i)$ consists of measured disturbance signal $forecasts$. Intuitively, a disturbance on the behavior change system on the time scale of interest is unlikely to be approximately represented as a ramp that significantly and continuously increases. Consequently, $f(q^{-1}, \alpha_d^{j_d})$ in these analyses similarly employs the Type-I structure in equation 3.91. By extension,

$$d_{flt}(k+i) = (1-\alpha_d)\, d(k+i) + \alpha_d d_{flt}(k+i-1) \tag{3.94}$$

Filtering the $Quit$ disturbance signal via an α_d^Q parameter would greatly affect performance since significant and immediate changes in both CPD and $Craving$ are induced by the $Quit$ step. However, detuning that slows the rejection of $Quit$ is operationally undesirable. Given that the $Quit$ step is the signal representing the fundamental reason why such an intervention is being designed, and is directly and inherently related to the CPD and $Craving$ target changes and corresponding dynamics, slowing the rate at which the full value of $Quit$ signal is delivered to

the controller is incongruous with the purpose of the intervention. Furthermore, filtering the *Quit* signal would also suggest that the anticipation of the *Quit* step should also be filtered, and the ability to forecast initiation of the upcoming quit attempt from the very first decision period is particularly advantageous; such *Quit* may facilitate aggressive pre-TQD dosing, which is more in line with current clinical practice. In other words, filtering *Quit* and providing an anticipated but filtered *Quit* signal unnecessarily diminishes the value of anticipating this measured disturbance. Filtering the *Quit* signal is also unnecessary from a practical point of view. As *Quit* initially brings CPD to its set point in the nominal case, filtering the *Quit* signal would lead to decision-making that essentially does not take advantage of the fact that *Quit* supports pursuit of the CPD target for a period of time, leading to more and unnecessary dosing around TQD. Due to these considerations, α_d^Q is omitted from Fig. 3.5 and $\alpha_d^Q = 0$ in the decision computations.

Anticipation of the *Stress* disturbance entails assuming that *Stress* will equal its most recent measured value for the duration of the prediction horizon, i.e., assumes $\{Stress(k+1), \cdots, Stress(k+p-1)\} = Stress(k)$. Consequently, calculation of the filtered *Stress* disturbance signal per equation 3.93 reflects $d(k+i) = Stress(k)$ for $0 \leq i \leq p-1$.

3.5.3 Unmeasured Disturbances

As $\mathcal{G}$ and b in equations 3.75 and 3.76 are both functions of $X(k)$ (see equations 3.79 and 3.82), an estimate for $X(k)$ must be calculated at each decision period. These estimates are obtained via a state observer. To be able to tune for measured disturbance rejection and for unmeasured disturbance rejection independently, a two step estimation procedure is employed:

1. *Estimation of* $X(k)$: An estimate of $X(k)$ is obtained per equations 3.95 and

3.96:

$$X(k|k-1) = \mathcal{A}X(k-1|k-1) + \mathcal{B}_1\Delta u(k-1) + \mathcal{B}_2\Delta\delta(k-1) + \mathcal{B}_3\Delta z(k-1) + \mathcal{B}_d\Delta d(k-1) \quad (3.95)$$

$$X(k|k) = X(k|k-1) + K_f(y(k) - \mathcal{C}X(k|k-1)). \quad (3.96)$$

Equation 3.95 is the prediction of $X(k)$ per the model (under the assumption of no white noise; see equation 3.35). Equation 3.96 involves a a correction when the effect of unmeasured disturbance is incorporated into the estimate, where the $(y(k) - \mathcal{C}X(k|k-1))$ term is the prediction error calculated with unfiltered d signals; K_f is a filter gain. Conceptually, K_f is a weight telling the controller how much to trust the model versus the measurement. Formulaically, K_f defines the speed and character of unmeasured disturbance rejection (Nandola and Rivera, 2013).

2. *Estimation of* $X_{flt}(k)$: $X_{flt}(k)$ represents an estimation of the augmented states using a filtered form of the measured disturbance signal:

$$X_{flt}(k|k-1) = \mathcal{A}X_{flt}(k-1|k-1) + \mathcal{B}_1\Delta u(k-1) + \mathcal{B}_2\Delta\delta(k-1) + \mathcal{B}_3\Delta z(k-1) + \mathcal{B}_d\Delta d_{flt}(k-1) \quad (3.97)$$

$$X_{flt}(k|k) = X_{flt}(k|k-1) + K_f(y(k) - \mathcal{C}X(k|k-1)). \quad (3.98)$$

The first term in equation 3.98 consists of the prediction obtained using the filtered measured disturbance while the second term is the prediction error, as before. The result is an estimate of the augmented states in which the effects of detuning the measured disturbance rejection via the α_d parameter(s) is decoupled from detuning the unmeasured disturbance rejection via K_f (Nandola and Rivera, 2013).

Ideally, an optimal K_f value would be determined. Doing so would require accurate matrices of the process noise and measurement noise covariances (Franklin *et al.*, 1998; Nandola and Rivera, 2013). However, these covariances are difficult to determine in this behavioral health problem setting. Instead, the filter gain is parameterized according to:

$$K_f = \begin{bmatrix} 0 \\ F_b \\ F_a \end{bmatrix} \tag{3.99}$$

where

$$F_a = \text{diag}\{(f_a) + 1, \ldots (f_a)_{n_y}\}, \;\; 0 < (f_a)_j \leq 1 \tag{3.100}$$

$$F_b = \text{diag}\{(f_b) + 1, \ldots (f_b)_{n_y}\} \tag{3.101}$$

$$(f_b)_j = \frac{(f_a)_j^2}{(1 + \alpha_j (f_a)_j)}, \;\; 1 \leq j \leq n_y. \tag{3.102}$$

$\alpha_j = 0$ or $= 1$ for rejection of step and ramp disturbances, respectively, in the j^{th} controlled variable. The speed of unmeasured disturbance rejection is proportional to value of the tuning parameter $(f_a)_j$; $(f_a)_j = 1$ corresponds to the no-detuning case, where very aggressive control action is taken to reject the influence of unmeasured disturbances on the j^{th} controlled variable (Lee and Yu, 1994; Nandola and Rivera, 2013).

CPD and $Craving$ are considered here to be the only signals that can be corrupted by unmeasured disturbances and the effects of plant-model mismatch, therefore, analysis of unmeasured disturbance rejection here considers f_a^{cpd} and f_a^{crav} parameters only.

Chapter 4

EVALUATING INTERVENTION PERFORMANCE

4.1 Overview

The objective of an HMPC-based smoking intervention is generally to promote successful cessation, and specifically to meet targets for the controlled variables; on this basis, performance requirements for the controller can be defined. These requirements would seek to balance the desire for intervention efficacy with a concern for dosing and resource-use demands. Ultimately, a clinician would determine what is acceptable performance. In this chapter, the intervention algorithm is assessed through simulation in two broad sets of analyses:

1. *Nominal Performance* — First, controller performance is evaluated when there is no model uncertainty and minimal plant-model mismatch (Morari and Zafiriou, 1989). In other words, the dose-, *Quit*-, and *Stress*-response models provided to the controller are those describing the particular patient who is the "plant" in the control loop. Nominal performance is first examined for an HMPC-framework that does not feature 3DoF tuning capabilities; similar analyses are later pursued for a formulation with 3DoF features. Nominal performance analysis in the following sections assumes nominal stability of the system (which can be observed via simulation).

2. *Robust Performance* — The controller is also evaluated for robust performance, i.e., performance in the presence of uncertainties (Morari and Zafiriou, 1989). Deriving formal robustness margins for constrained predictive control is a significant undertaking (Rossiter, 2003) and lies outside the scope of this work. In

this chapter, uncertainties in the model are simulated via model perturbations. Initially, unmodeled exogenous disturbances are incorporated in the form of measurement noise. Plant-model mismatch is then incorporated in the form of different parametric uncertainties in the open-loop models of the representative patient. Next, the intervention is evaluated when the plant consists of models estimated for single subjects from the placebo group in the McCarthy *et al.* (2008b) study who are not the subject upon which the representative patient is based. The assessment of robust performance includes robust stability.

4.2 Evaluating the Intervention: Time Frame & Conditions of Interest

The potential clinical utility of an HMPC-based adaptive smoking cessation intervention is illustrated through simulation. Simulink, MATLAB's graphical simulation environment (The MathWorks, 2014g), facilitates such performance, robustness, and tuning analyses, and a clinician-friendly tuning strategy is ultimately defined. These analyses focus on the following conditions.

4.2.1 Time Frame

Intervention performance is primarily examined for a 50 day time span: control action is first determined on day 0 and TQD is day 15.

Implementing the intervention approximately two weeks pre-TQD reflects the recommendation that patients pick a TQD two weeks prior to initiation of the actual quit attempt (Tobacco Use and Dependence Guideline Panel, 2008). Calculating potential control actions for days or weeks pre-TQD offers a number of opportunities. If suitable, nonzero dosage levels can be assigned in anticipation of TQD. Fixed intervention protocols that feature such pre-TQD dosing have been shown to be helpful for treatments featuring counseling and bupropion (Tobacco Use and Dependence

Guideline Panel, 2008; McCarthy *et al.*, 2008b; U.S. Department of Health and Human Services and Centers for Disease Control and Prevention, 2014; Lexicomp, 2014) and is being explored for fixed interventions employing nicotine replacement therapies (Bullen *et al.*, 2006; Carpenter *et al.*, 2013).

After TQD, analysis focuses on the first five weeks of the quit attempt. Approximately 70% of quit attempts fail within one month of initiation of the attempt (Meyer *et al.*, 2003), suggesting intervention evaluation on this time scale is appropriate. Furthermore, habituation is more probable over long time periods (e.g., where a lozenge has differential effects if delivered one week vs. ten weeks post-TQD; Martín *et al.*, 2014). It is also hypothesized that the overall long-term cessation process involves discrete stages characterized by distinct individual processes (DiClemente *et al.*, 1991; Prochaska and DiClemente, 1983; Velicer *et al.*, 1998). These factors suggest it may not be reasonable to assume that the models describing dynamics over the weeks following TQD, such as those in equations 3.5 through 3.9 and Table 3.2, will be representative of the dynamics that would be observed on the order of months post-TQD. Consequently, a quit attempt is defined here as having been unsuccessful if a patient has not significantly reduced smoking levels by day 50.

4.2.2 Disturbances

The *Quit* disturbance is described by equation 3.11 with TQD = 15, unless otherwise specified.

Stress also acts as a measurable and anticipatable disturbance. A patient can experience stress on a variety of time scales (Ehlert and Straub, 1998; Shiffman and Waters, 2004), impacting smoking behavior through a number of pathways (Shiffman and Waters, 2004). In the following, *Stress* is represented in two ways—as a step of magnitude 3 and as stochastic signal. These signal types are intended to reflect

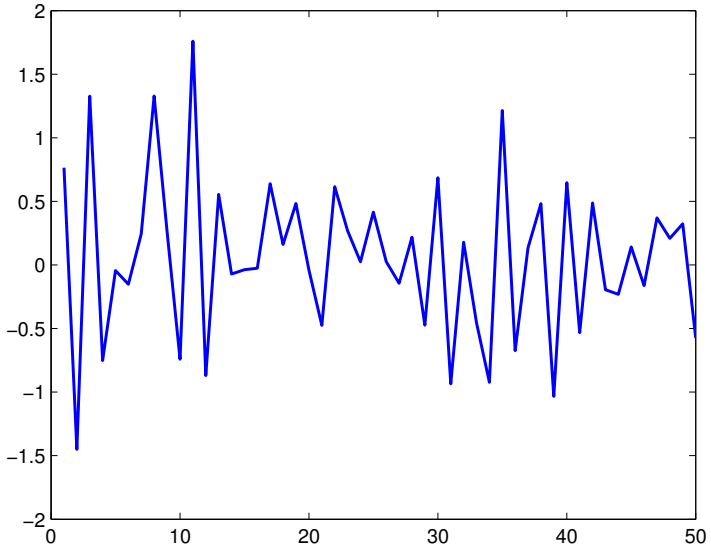

Figure 4.1: A Time Series Realization of the Stochastic Signal that Represents *Stress* (see Equation 4.1).

that stressors and their effects may be acute or chronic in nature (Ehlert and Straub, 1998; Shiffman and Waters, 2004). Day-to-day changes in the stochastic *Stress* signal are assumed to be autocorrelated (DeLongis *et al.*, 1988). Specifically, the stochastic *Stress* signal is represented by a normally distributed pseudorandom sequence ($N \sim (0, 0.5)$) filtered according to the first order autoregressive model in equation 4.1 (The MathWorks, 2014b),

$$Stress(t) = \frac{1.67}{(0.67\, q^{-1} + 1)}\, a(t) \tag{4.1}$$

where $Stress(t)$ and $a(t)$ are discrete signals and q^{-1} is the time-domain backshift operator. Fig. 4.1 features a realization of this stochastic stress disturbance.

4.3 Performance Metrics

Visual examination of simulated CPD and $Craving$ responses and the corresponding dosing profiles facilitate assessment of the intervention formulation and different tuning conditions. However, quantitative metrics that are used here to help evaluate

nominal and robust performance include:

- Rise time, t_r^y – Length of time after TQD required for the response of controlled variable y to first reach within 5% of its target level.
- Offset, $e_{t_i}^y$ – Deviation between controlled variable y and its target at time t_i (in units of y).
- *Total Cigs* – Total number of cigarettes smoked over a specified time interval (generally for $t \geq$ TQD, below).
- $Tot(u_*)$ – Total number of doses of intervention component u_* assigned over a specified time interval.
- Mean, μ^x – Average value of x over the time interval of interest; x may represent a controlled or manipulated variable.
- Maximum value, $max(x)$ – Maximum value of the sequence of variable x between two specified time points; x may represent a controlled or manipulated variable.
- Minimum value, $min(x)$ – Minimum value of the sequence of variable x between two specified time points; x may represent a controlled or manipulated variable.
- Variance, $var(x)$ – Variance of the sequence of variable x between two specified time points; x may represent a controlled or manipulated variable.
- Cumulative deviation from target, J_e^y – Metric reflecting the cumulative deviation between controlled variable y and its target level between times t_1 and t_2; calculated according to (Nandola and Rivera, 2013):

$$J_e^y = \sum_{k=t_1}^{t_2} (y(k) - y^{target}(k))^T (y(k) - y^{target}(k)) \tag{4.2}$$

where $y^{target}(k)$ is the user-defined target for y at time k (which is equivalent to the reference trajectory provided to the controller in the no-3DoF and $\alpha_r^y = 0$ cases). This metric is intended to quantify total deviation of a controlled variable from its target; as the simulations examined in this section enforce saturation in CPD and $Craving$ levels such that these controlled variables can never assume negative values, the values of J_e^y and of $Tot(y)$ are directly related.

- Cumulative intervention energy, $J_I^{u_*}$ – Metric reflecting the cumulative energy required of manipulated variable u_* between times t_1 and t_2; calculated according to (Nandola and Rivera, 2013):

$$J_I^u = \sum_{k=t_1}^{t_2} (u(k))^T (u(k)) \tag{4.3}$$

- Total relaxation required, ϵ_T – The total amount of constraint relaxation required for a given scenario and specified Q_ϵ, defined as the sum of all values of the slack calculated at each decision period, summed across all decision periods. ϵ_T should offer insight into the amount of relaxation necessary to obtain feasible solutions to the optimization computations associated with each control move.

While these metrics offer quantified insight into intervention performance, each one is not examined for each scenario below, and most of the following discussion focuses on general characteristics of performance.

4.4 Nominal Performance

The goal of the simulations that follow in this section is to identify how the penalty weights can be adjusted to give favorable CPD and $Craving$ performance and desirable dosing schedules. Unless otherwise specified, the following parameters are kept constant: $p = 30$ days, $m = 7$ days, $Q_{cpd} = 10$, $Q_{\Delta u_*} = Q_{u_*} = Q_\delta = Q_z = 0$

(where $*$ indicates counseling, bupropion, or lozenge here), $Q_\epsilon = 1e5$, and $T = 1$ day. To avoid unnecessary dosing, WIP_T (see equation 3.19) is assigned a set point of zero for all time points. Initially, the bupropion move size constraint in equation 3.58 is implemented, which allows decreases in u_b over time, where $T_{sw} = 4$ and $t^1_{sw} = 2$. In several cases it is assumed $Stress$ does not fluctuate beyond some baseline level, mathematically represented as $Stress = 0$.

The following sections incorporate increasing degrees of penalty weighting and tuning complexity. The scenarios of focus for nominal performance analysis are described in Table 4.1 and referred to by scenario number in the remainder of Section 4.4.

Table 4.1: Scenarios Considered in Section 4.4.

Scenario	Q_{wip_T}	α^*_r	α^S_d	$Stress$	Δu_b constraint equation
1	0	0	0	0	3.57
2	≥ 0	0	0	0	3.57
3	≥ 0	0	0	Step	3.57
4	≥ 0	0	0	Stochastic	3.57
5	≥ 0	0	0	Stochastic	4.4
6	1	≥ 0	0	0	4.4
7	1	0	≥ 0	Step	4.4
8	1	0	≥ 0	Stochastic	4.4
9	1	≥ 0	≥ 0	Stochastic	4.4

4.4.1 Tuning via Objective Function Penalty Weights

Within a standard HMPC framework (Fig. 3.1), the character of the control action and controlled variable responses is determined by the objective function penalty weights for a given set of targets, constraints, and manipulated variables. In this con-

text, Q_{cpd} and Q_{crav}, penalty weights on the CPD and $Craving$ errors, respectively, directly promote dosing that pursues intervention targets. Q_{wip_T} penalizes deviations of WIP_T from its set point, which when set to 0 reflects a concern for the total dosing demands for a given day. Q_{u_*} penalizes deviations of manipulated variable u_* (where u_* denotes u_c, u_b, or u_l) from its set point. $Q_{\Delta u_*}$ can suppress dosing changes. Altering the values of these penalty weights relative to one another can facilitate a practical desire to prioritize intervention goals, influence the time-varying nature of controlled variable offsets, and adjust the character of dosing schedules.

Scenario 1

If the only nonzero objective function penalty weights are those for CPD and $Craving$ set point tracking, adjustment of the manipulated variables is aggressive, limited only by the availability of discrete-valued dosage levels and hard constraints. In other words, with $Q_{wip_T} = Q_{u_*} = Q_{\Delta_{u_*}} = 0$, dosing for the purposes of smoking abstinence and $Craving$ reduction tracking will be aggressive.

Mathematically, the relative values of Q_{cpd} and Q_{crav} determine which outcome target is prioritized during day-to-day decision making. Fig. 4.2 depicts three sets of responses where $Q_{cpd} = 10$ and Q_{crav} is an order of magnitude smaller than, equal to, and an order of magnitude larger than Q_{cpd} (Fig. 4.2a, Fig. 4.2b, and Fig. 4.2c, respectively). The performance metrics corresponding to these figures (as well as $Q_{crav} = \{5, 20, 50\}$), are found in Table 4.2.

As seen in Fig. 4.2, various values of Q_{crav} relative to Q_{cpd} lead to CPD and $Craving$ responses with relatively good set point tracking and similar general features. The patient is able to quit relatively successfully, reflected by the only minor degrees of lapse, which occur according to various trends. Each penalty weight combination features some CPD offset at day 50, although this offset appears to be due to lapse

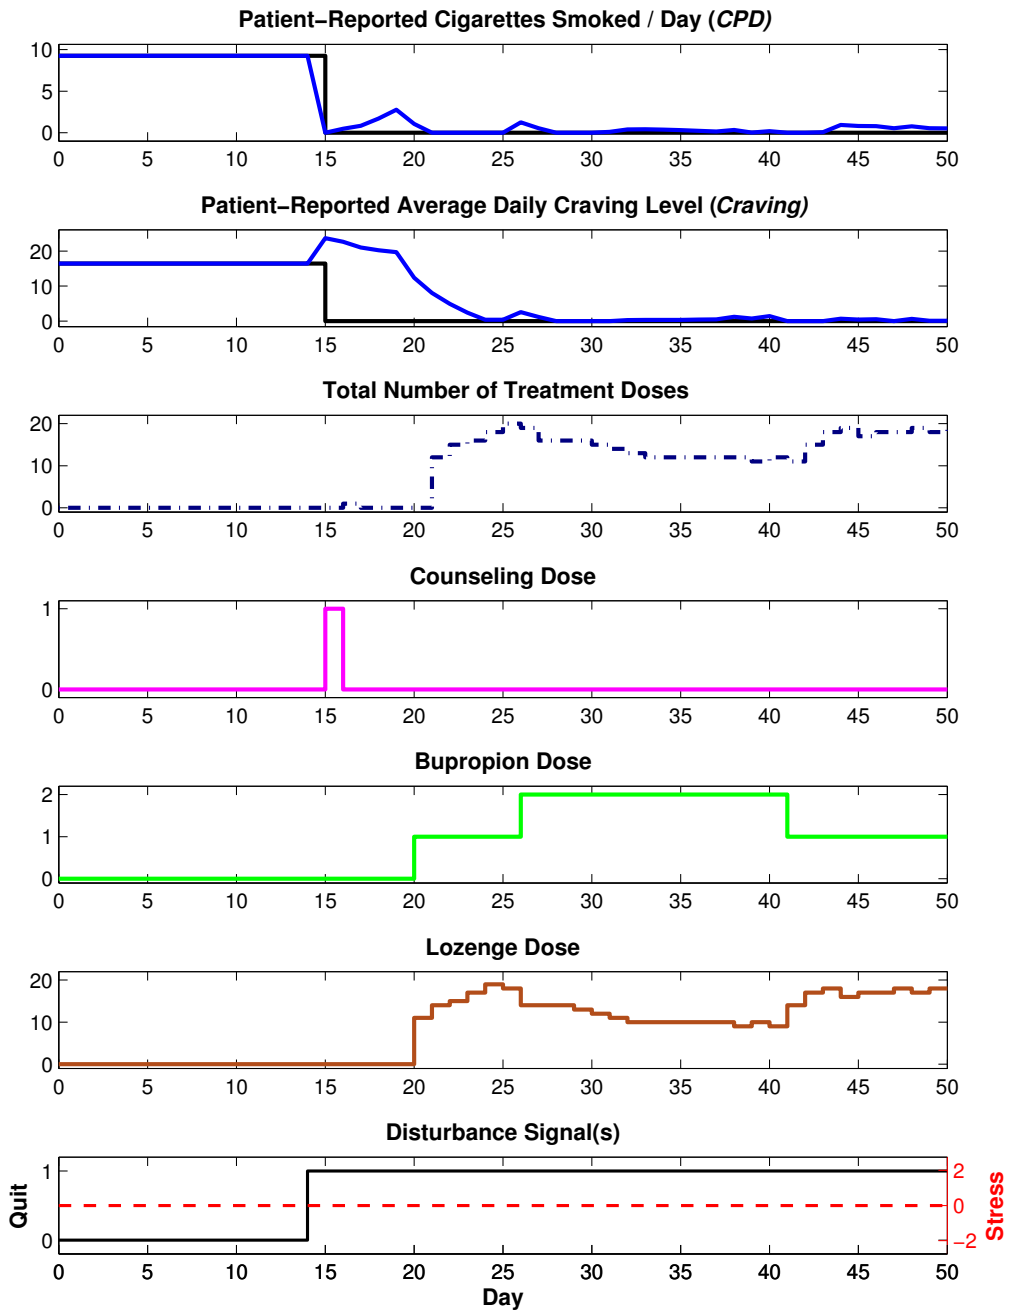

(a) Scenario 1: Outcome and Manipulated Variable Responses Where $Q_{wip_T} = 0$, $Q_{cpd} = 10$, and $Q_{crav} = 1$.

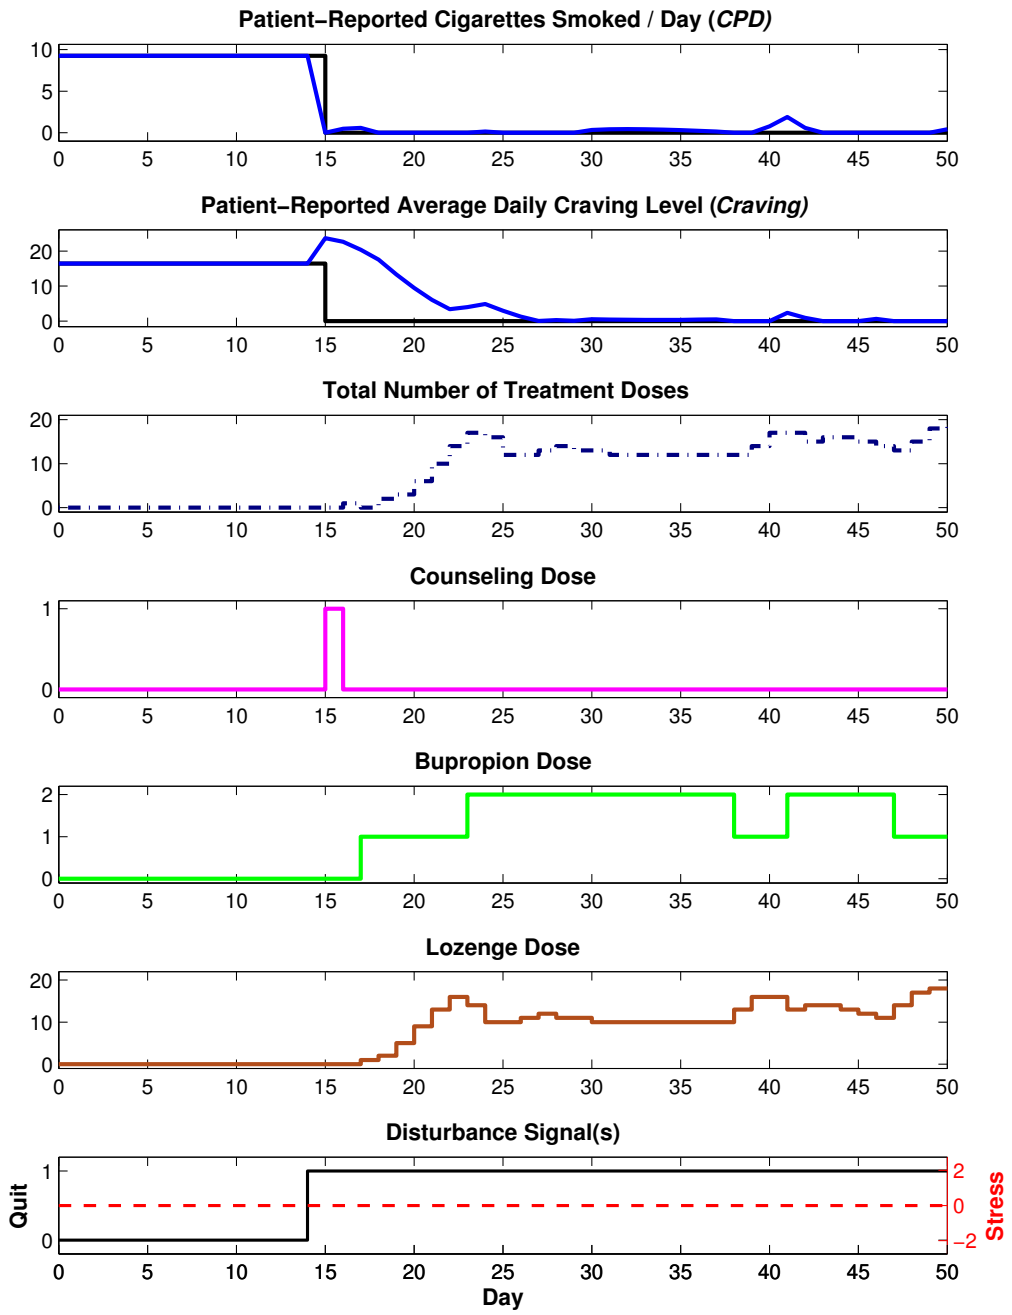

(b) Scenario 1: Outcome and Manipulated Variable Responses where $Q_{wip_T} = 0$, $Q_{cpd} = 10$, and $Q_{crav} = 10$.

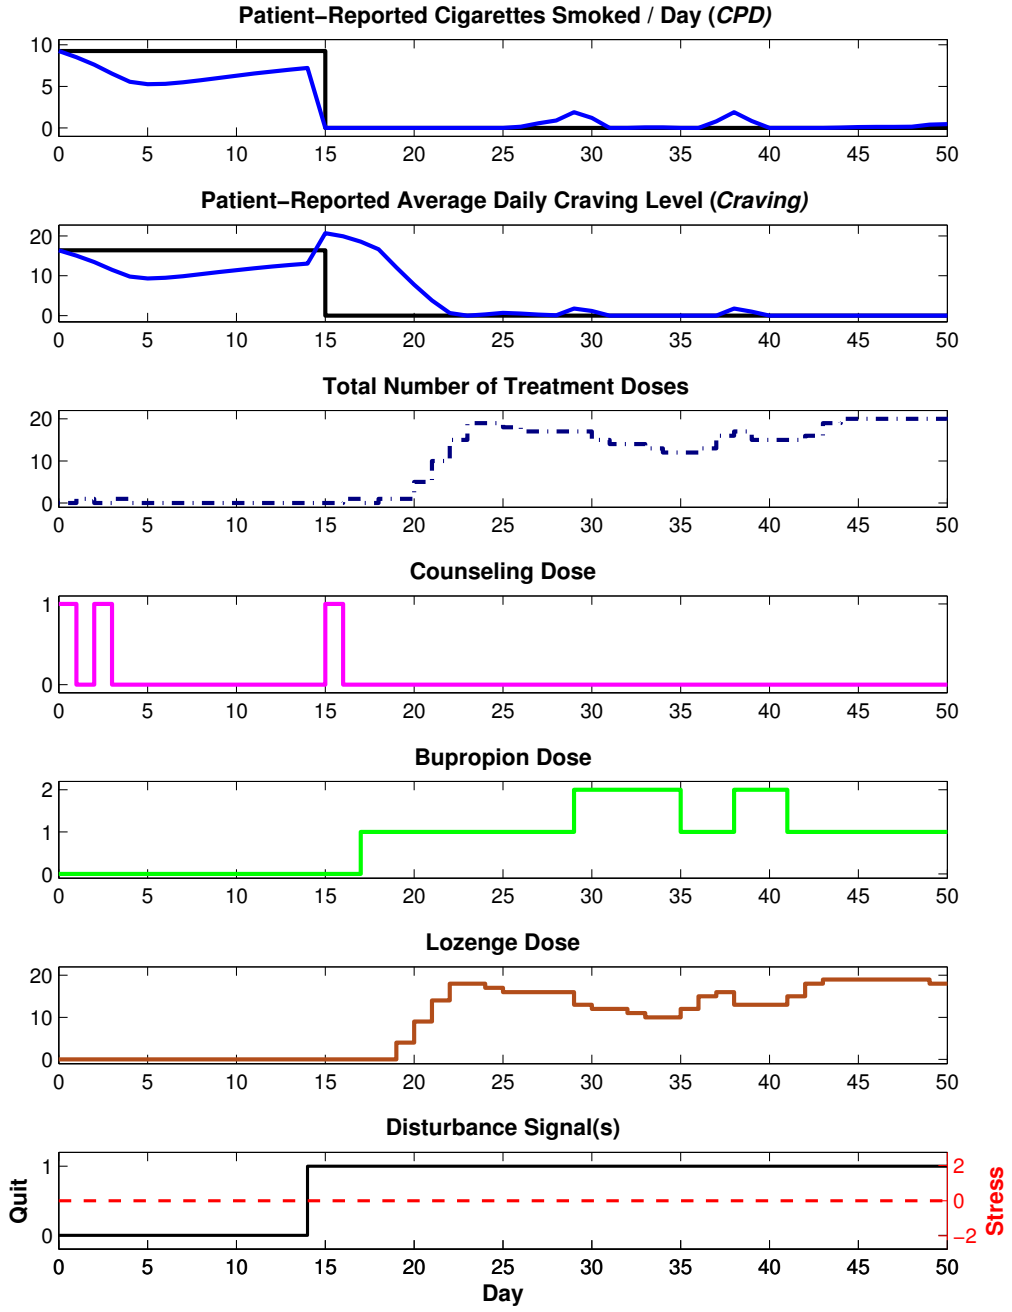

(c) Scenario 1: Outcome and Manipulated Variable Responses where $Q_{wip_T} = 0$, $Q_{cpd} = 10$, and $Q_{crav} = 100$.

Figure 4.2: Scenario 1: Outcome and Manipulated Variable Responses where $Q_{wip_T} = 0$, $Q_{cpd} = 10$ for Various Q_{crav} values.

Table 4.2: Scenario 1: Performance Metrics for the Intervention with Various Penalty Weight Values ($Stress = 0$).

	Figure:	4.2a		4.2b			4.2c
Parameters	Q_{cpd}	10	10	10	10	10	10
	Q_{crav}	1	5	10	20	50	100
	Q_{wip_T}	0	0	0	0	0	0
Performance, $t \geq$ TQD	t_r^{cpd}	0	0	0	0	0	0
	t_r^{crav}	9	11	12	12	12	7
	e_{50}^{cpd}	0.53	0.78	0.42	0.53	0.53	0.46
	e_{50}^{crav}	0.10	0.72	0.00	0.10	0.10	0.00
	Total Cigs	16.33	16.94	7.65	11.95	17.66	9.86
	Days CPD=0	11	11	20	14	14	18
	$max(cpd)$	2.77	2.57	1.90	0.89	1.30	1.90
	$max(crav)$	23.70	23.70	23.70	23.70	23.70	20.67
	$max(WIP_T)$	20	19	19	19	19	20
	J_e^{cpd}	19.06	21.58	6.34	7.65	16.32	11.62
	J_e^{crav}	2574.94	2202.51	2173.27	2167.42	2166.47	1674.79
	$J_e^{wip_T}$	7212	7498	6045	6833	7913	8396
	$var(cpd)$	0.33	0.39	0.13	0.11	0.22	0.25
	$var(crav)$	56.21	47.38	47.09	47.23	47.69	38.67
	$var(WIP_T)$	39.58	34.41	25.28	30.92	34.75	39.88
	$J_I^{u_c}$	1	1	1	1	1	1
	$J_I^{u_b}$	76	70	97	79	61	61
	$J_I^{u_l}$	6304	6593	5048	5948	7082	7527
	$var(u_c)$	0.03	0.03	0.03	0.03	0.03	0.03
	$var(u_b)$	0.49	0.32	0.37	0.35	0.28	0.28
	$var(u_l)$	33.37	29.79	21.03	27.26	30.00	35.48
	ϵ_T	0.28	1.48	1.51	3.37	5.75	35.33

events occurring just prior to day 50 specifically, after substantial periods of no or negligible smoking.

A scenario with Q_{cpd} non-trivially greater than Q_{crav} intuitively seems like it should provide optimal performance in terms of supporting cessation success most directly. However, post-TQD deviations in CPD, $Craving$, and WIP_T from target levels appear to be greater when $Q_{cpd} \gg Q_{crav}$ for this nominal case; Fig. 4.3 indicates that cumulative post-TQD outcome errors approach minimal levels when Q_{cpd} and Q_{crav} are on the same order of magnitude. Specifically, these plots suggest that the total amount of post-TQD smoking does not significantly improve when Q_{crav} increases in value beyond Q_{cpd}; $Q_{cpd} = Q_{crav} = 10$ also corresponds to low relative values of both J_e^{crav} and $J_e^{wip_T}$, and larger Q_{crav} values similarly do not systematically improve the performance in both metrics.

Examining outcome variable responses and dosing schedules specifically, Fig. 4.2 suggest that when both weights equal 10, CPD features very minor instances of lapse during the first weeks of quitting and only a few subsequent days where more than one whole cigarette is smoked. Furthermore, the magnitudes of the total amount of smoking during the quit attempt ($Total\ Cigs$, also reflected in J_e^{cpd}) and CPD offset on day 50 are smallest for the equal weight case. Similarly, the total number of days the patient does not smoke for this weight combination, 20 days, is nearly double that for the $Q_{crav} \ll Q_{cpd}$ cases and more than 25% greater than the $Q_{crav} = 20$ and 50 cases. This, and the low variance in the CPD sequence post-TQD, suggest that significantly disparate penalty weights on the CPD and $Craving$ set point tracking goals can lead to performance degradation.

$Craving$ consistently increases above baseline levels around TQD despite the intervention. The similarity of the $Craving$ inverse response, t_r^{crav} values, and $max(crav)$ values across the various penalty weight combinations are likely a consequence of the

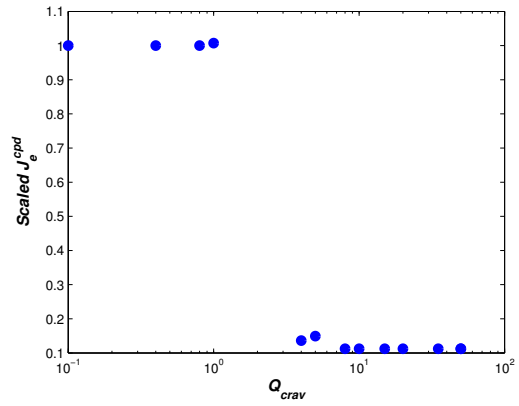

(a) Relative J_e^{cpd} Values.

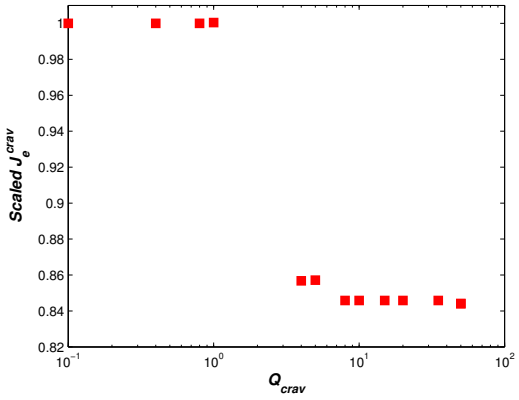

(b) Relative J_e^{crav} Values.

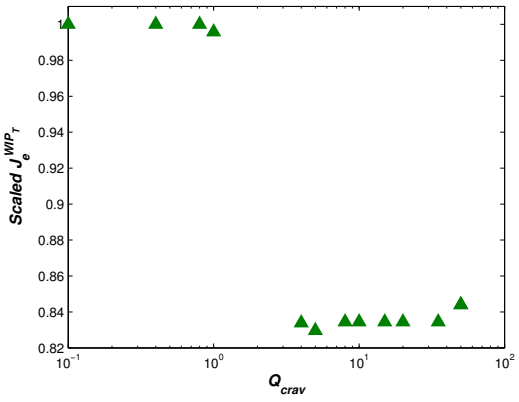

(c) Relative $J_e^{wip_T}$ Values.

Figure 4.3: Scenario 1: Metrics of Post-TQD Intervention Performance for $Q_{wip_T} = 0$, $Q_{cpd} = 10$, and Varying Levels of Q_{crav}, Relative to the $Q_{crav} = 0.1$ Case.

lower bound on CPD (see equation 3.59) and the interrelationship between CPD and $Craving$. As $Quit$ drives CPD to its lower limit—and equivalently, its set point—the controller avoids dosing around TQD that would drive CPD below its lower bound, even though such control action would also bring $max(crav)$. Fig. 4.2c supports this assertion: here, the peak in $Craving$ on TQD, 20.67 points, is lower than the other scenarios, all 23.70 points; however, the value of ϵ_T for this case relative to the equivalent simulations suggests that to achieve the lower $Craving$ maximum as well as feasible solutions for the $Q_{crav} = 100$ problem, significant relaxation of the hard constraint bounds is required—likely corresponding to significant relaxation of physically realizable CPD level.

Fig. 4.2 suggests that the various combinations of penalty weights generally assign unit counseling around TQD and rely heavily on bupropion and lozenge dosing to achieve set point tracking, often trading off between use of lozenges and implementation of a second bupropion dose as time goes on. The dosing around TQD is relatively limited, with u_b and u_l being gradually assigned until becoming very aggressive approximately one week after TQD (for the cases with limited constraint relaxation). These u_b moves contrast current clinical practice, which prescribe the maximum bupropion dose four days *prior* to TQD (Tobacco Use and Dependence Guideline Panel, 2008; Lexicomp, 2014). This relative delay in aggressive dosing reflects the fact that the nominal patient is initially above to achieve the $CPD = 0$ target without the aid of any treatment components; the manipulated variables are then assigned in a manner that addresses the resumption of smoking—gradual departure from target CPD levels—that would otherwise occur.

In terms of overall dosing requirements for this favorable penalty weight combination, the intervention relies on greater magnitudes of and greater adjustments to lozenge dosage and less on bupropion dosage when $Q_{cpd} \neq Q_{crav} \neq 10$ (see $J_I^{u_b}$, $J_I^{u_l}$,

and $var(u_l)$ in Table 4.2). Specifically, Fig. 4.2b employs an initial dose of bupropion just after TQD and employs the second bupropion dose more than the other cases in Fig. 4.2, which also corresponds to lower peak and steady-state u_l levels in the first three weeks of the quit attempt.

Scenario 2

As alluded to previously, assigning $Q_{u_*} > 0$ values and $u_*^{ref} = 0$ would serve to discourage assignment of nonzero treatment dosages, and assigning $Q_{\Delta u_*} > 0$ for move suppression would serve to discourage changes of these manipulated variables. Mathematically, this approach treats each manipulated variable separately and a concern for total dosing or over-dosing would require careful selection of synergistic combinations of Q_{cpd}, Q_{crav}, Q_{u_c}, Q_{u_b}, Q_{u_l}, $Q_{\Delta u_c}$, $Q_{\Delta u_b}$, and $Q_{\Delta u_l}$ values. Relatedly, the concern is primarily in managing dosing in aggregate, and managing such a concern via individual manipulated variable penalty weights restricts dosing flexibility in an unintuitive manner and to an unnecessary degree. Instead, a construct accounting for total dosing over a given period of time is defined and given a set point.

Fig. 4.4 features J_e^* values (where $*$ indicates CPD or $Craving$) corresponding to various Q_{wip_T} levels and $Q_{cpd} = Q_{crav} = 10$, scaled relative to the corresponding J_e^* values where $Q_{wip_T} = 0$. Fig. 4.5 features controlled variable responses and the corresponding dosing schedules when Q_{wip_T} assumes values orders of magnitude smaller than or equal to $Q_{cpd} = Q_{crav} = 10$. The corresponding performance metrics are found in Table 4.3.

These figures suggest that increasing $Q_{wip_T} > 1$ in this scenario quickly degrades CPD performance significantly, as well as $Craving$ performance. This could be expected, as moving WIP_T toward its set point of no dosing becomes a greater priority relative to the CPD and $Craving$ targets, which leads to less dosing—a

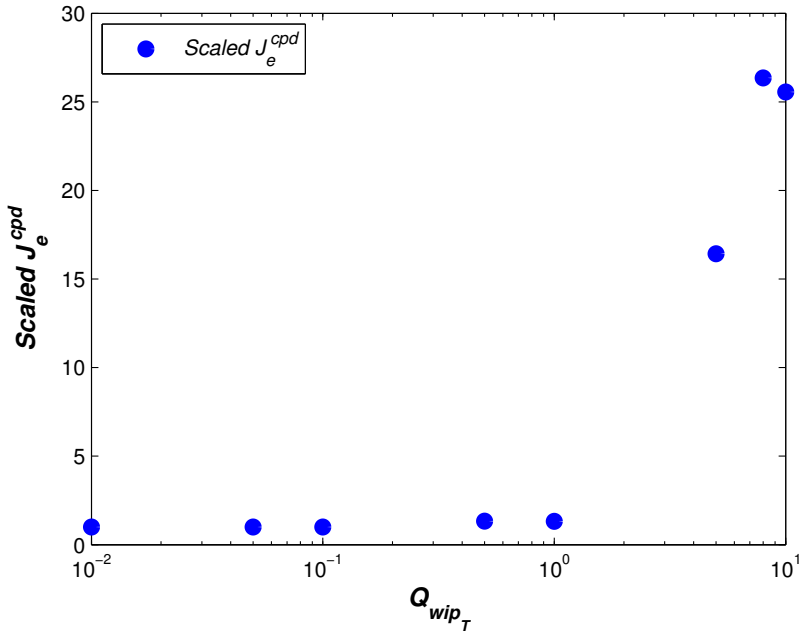

(a) Relative J_e^{cpd} Values.

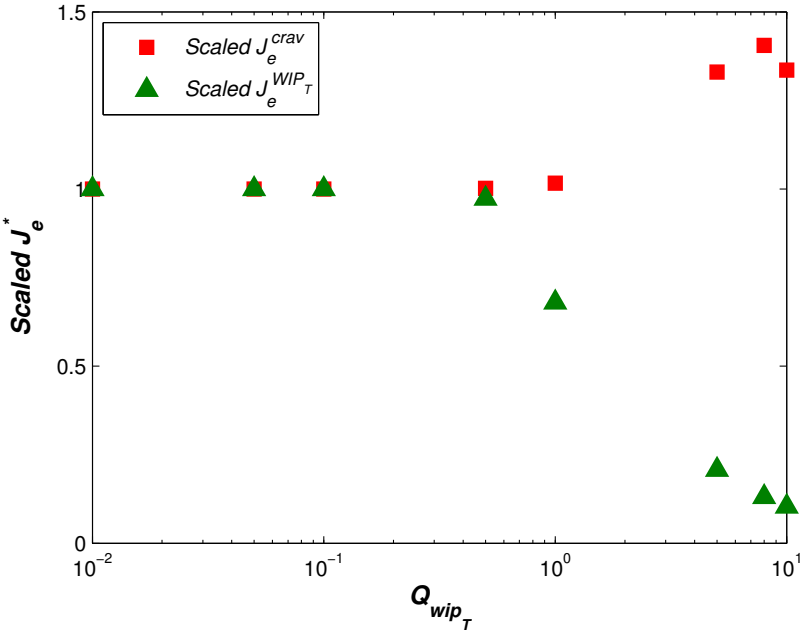

(b) Relative J_e^{crav} (Red squares) and $J_e^{wip_T}$ (Green Triangles) Values.

Figure 4.4: Scenario 2: Metrics of Post-TQD Intervention Performance for Varying Q_{wip_T} Values and $Q_{cpd} = Q_{crav} = 10$, Scaled Relative to the $Q_{wip_T} = 0$ Case.

direct consequence of which is increased post-TQD smoking and craving levels. This is particularly apparent in Fig. 4.5c. In this case, the maximum daily lozenge dose is 6 lozenges, and only 41 u_b doses are assigned, versus more than 50 for the $Q_{wip_T} = 0.1$ and $Q_{wip_T} = 1$ cases. Such dosing suppression leads to significantly worse performance compared to the other cases in Table 4.3.

Compared to the scenario where $Q_{wip_T} = 0$ (Fig. 4.2b), the $Q_{wip_T} = 0.1$ and $Q_{wip_T} = 1$ cases feature better CPD performance according to some metrics. Notably, the highest $max(cpd)$, J_e^{cpd}, and J_e^{crav} magnitudes are smaller. That said, there are six fewer cigarette-free days with these two penalty values. In terms of dosing, a WIP_T penalty weight two orders of magnitude lower than the CPD and $Craving$ penalty weights does not sufficiently smooth the u_l dosing profile, as quantified through the $var(u_l)$ value, and still requires high amounts of lozenge consistently. Interestingly, a Q_{wip_T} value that is approximately one order of magnitude smaller than or equal to Q_{cpd} and Q_{crav} appears to fall within a window of Q_{wip_T} values that decreases the total dosing demands by approximately 25% compared to smaller magnitudes of the WIP_T penalty weight, yet negatively affect total post-TQD CPD and $Craving$ set point deviations only marginally. The $Q_{wip_T} = 1$ case appears to retain the suppressed total dosing and less variant u_l adjustments while still achieving greater CPD set point tracking during the quit attempt. These features are achieved simultaneously by the controller's reliance on additional counseling and maximum bupropion dosages as compared to the $Q_{wip_T} = 0$ case. Fig. 4.2a also highlights that with $Q_{wip_T} = 1$ here, the intervention relies more consistently on u_b doses, which is more in congruent with current clinical practice (Lexicomp, 2014).

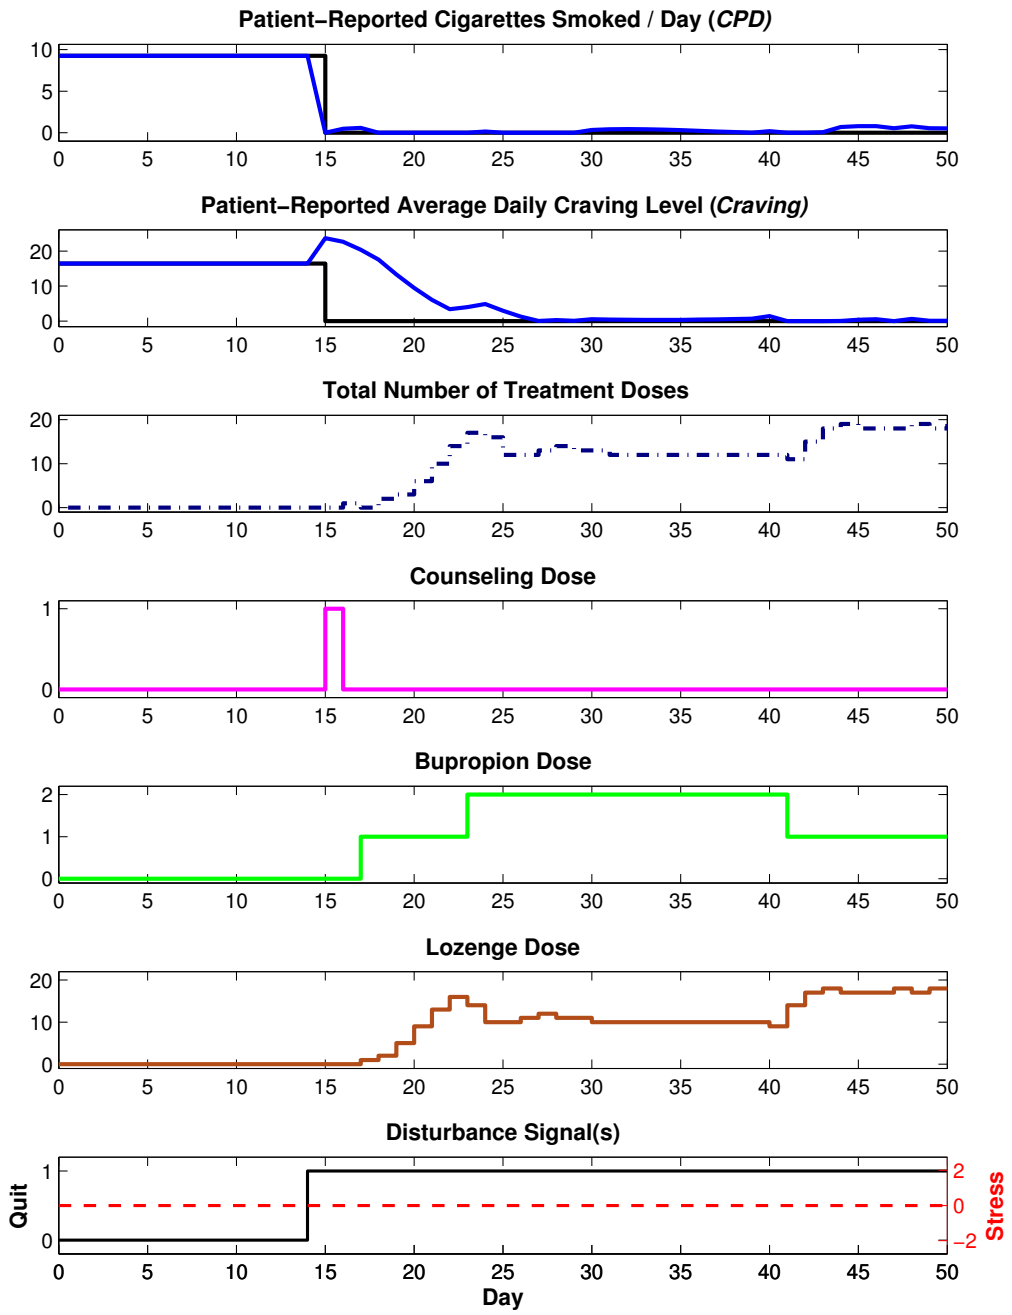

(a) Scenario 2: Outcome and Manipulated Variable Responses where $Q_{cpd} = Q_{crav} = 10$ and $Q_{wip_T} = 0.1$.

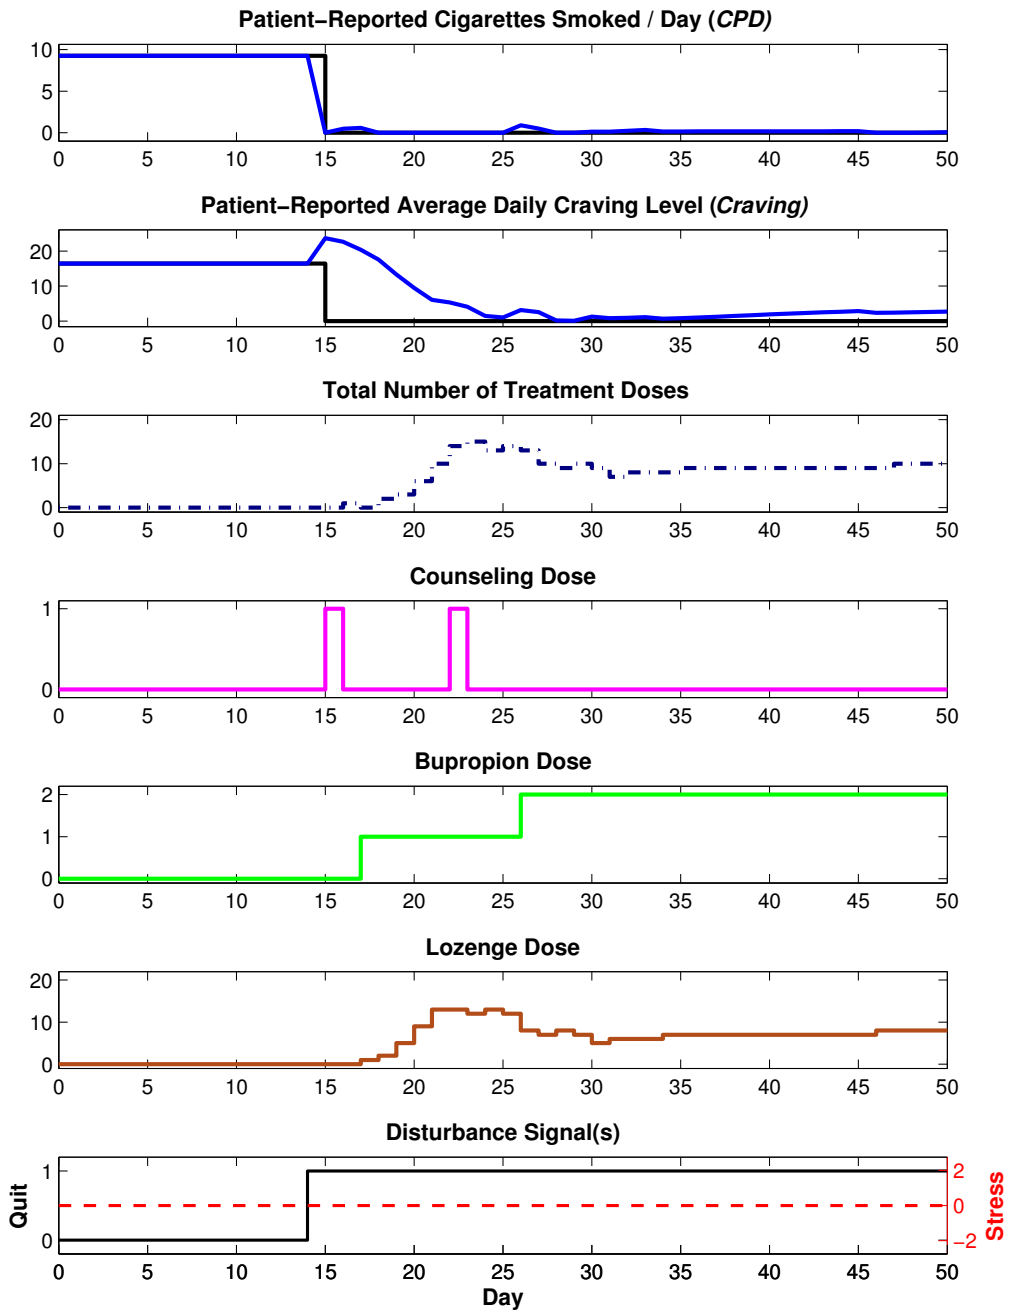

(b) Scenario 2: Outcome and Manipulated Variable Responses where $Q_{cpd} = Q_{crav} = 10$ and $Q_{wip_T} = 1$.

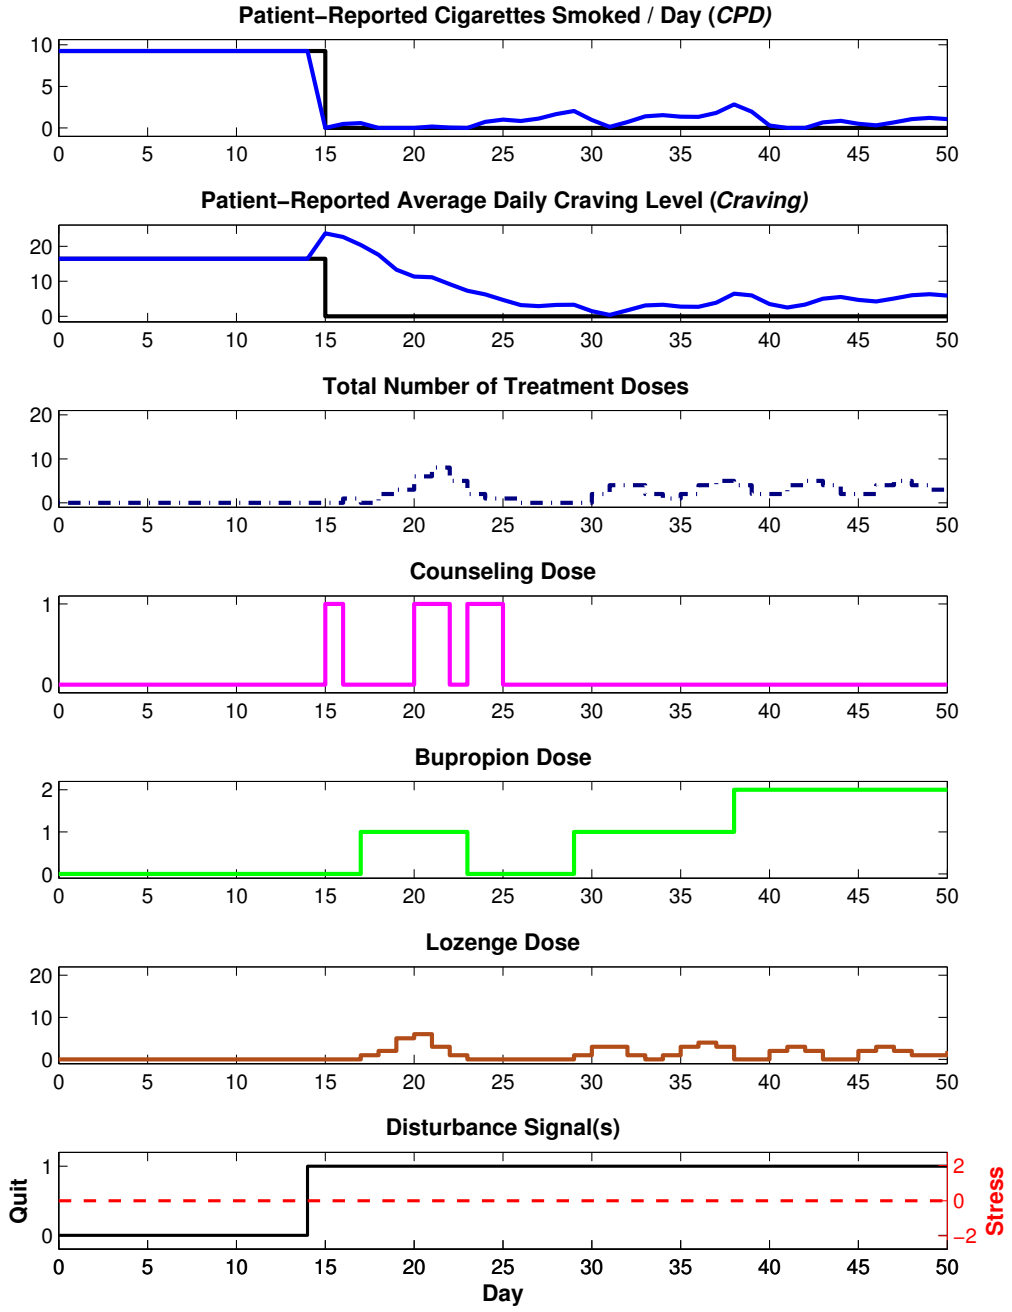

(c) Scenario 2: Outcome and Manipulated Variable Responses where $Q_{cpd} = Q_{crav} = 10$ and $Q_{wip_T} = 10$.

Figure 4.5: Scenario 2: Outcome and Manipulated Variable Responses where $Q_{cpd} = Q_{crav} = 10$ and Various Q_{wip_T} Values.

Table 4.3: Scenario 2: Performance Metrics for the Intervention with Various Penalty Weight Values, $Q_{wip_T} \neq 0$ ($Stress = 0$).

	Figure:	4.5a	4.5b		4.5c
Parameters	Q_{cpd}	10	10	10	10
	Q_{crav}	10	10	10	10
	Q_{wip_T}	0.1	1	5	10
Performance, $t \geq$ TQD	t_r^{cpd}	0	0	0	0
	t_r^{crav}	12	13	17	16
	e_{50}^{cpd}	0.53	0.06	0.19	1.07
	e_{50}^{crav}	0.10	2.72	4.20	5.91
	Total Cigs	9.02	5.55	22.70	29.52
	Days CPD=0	14	14	10	6
	$max(cpd)$	0.80	0.91	2.15	2.84
	$max(crav)$	23.70	23.70	23.70	23.70
	$max(WIP_T)$	19	15	7	8
	J_e^{cpd}	4.90	2.26	28.21	41.66
	J_e^{crav}	2170.04	2248.06	2822.83	2864.20
	$J_e^{wip_T}$	35	119	576	693
	$var(cpd)$	0.08	0.04	0.40	0.50
	$var(crav)$	46.87	41.48	41.70	34.50
	$var(WIP_T)$	0.03	0.25	2.21	2.65
	$J_I^{u_c}$	1	2	5	5
	$J_I^{u_b}$	88	109	61	67
	$J_I^{u_l}$	5457	2173	404	171
	$var(u_c)$	0.03	0.05	0.12	0.12
	$var(u_b)$	0.37	0.35	0.41	0.58
	$var(u_l)$	25.74	9.67	3.61	2.48
	ϵ_T	2.28	1.49	2.46	1.81

Scenario 3

Nominal performance is first examined where $Stress$ beyond some baseline level takes the form of a step change of magnitude three. Given the risk for relapse posed by a disturbance of this nature, performance is evaluated when the step disturbance occurs two weeks into the quit attempt, i.e., $Stress(t) = 0, t < 24$ and $= 3$, $t \geq 29$. Fig. 4.6 features plots with $Q_{cpd} = 10$ and combinations of Q_{crav} and Q_{wip_T}. The corresponding performance metrics are found in Table 4.4.

Poor performance results when Q_{wip_T} is on the same order of magnitude as that of the penalty weight(s) for CPD and/or $Craving$ set point tracking. Specifically, the control action for the $Q_{crav} = Q_{wip_T} = 1$ (Fig. 4.6b) and $Q_{cpd} = Q_{crav} = Q_{wip_T} = 10$ cases is effectively unable to reject the $Stress$ disturbance, where approximately 3 and 5 cigarettes are smoked on day 50, respectively. Such poor performance is confirmed by similarly high $Craving$ levels at day 50, large *Total Cigs*, $max(cpd)$, $max(crav)$, J_e^{cpd}, and J_e^{crav} values, and relatively small number of cigarette-free days for these cases. Conversely, the scenarios where $Q_{wipT} = 0$ feature effective disturbance rejection; e.g., $e_{50}^{cpd} = 0.00$ and $e_{50}^{crav} < 1$ for Fig. 4.6a. That said, even aggressive dosing—where u_l dosing is sustained at maximum levels—features post-TQD days where more than 2 cigarettes are smoked.

The $Q_{cpd} = Q_{crav} = 10$ with $Q_{wip_T} = 1$ case appears to balance step disturbance rejection and peak daily u_l requirements effectively in the nominal performance case. While it does feature consecutive days of post-TQD smoking, there are 17 cigarette-free days (the greatest number of smoke-free days of the comparable cases in Table 4.4) and the maximum lapse is only 1.85 cigarettes. As a whole, these tuning weights together offer better overall set point tracking despite the step $Stress$ disturbance, as compared to the $Q_{wip_T} = 0$ case shown. This is quantified by J_e^{cpd}, which is nearly

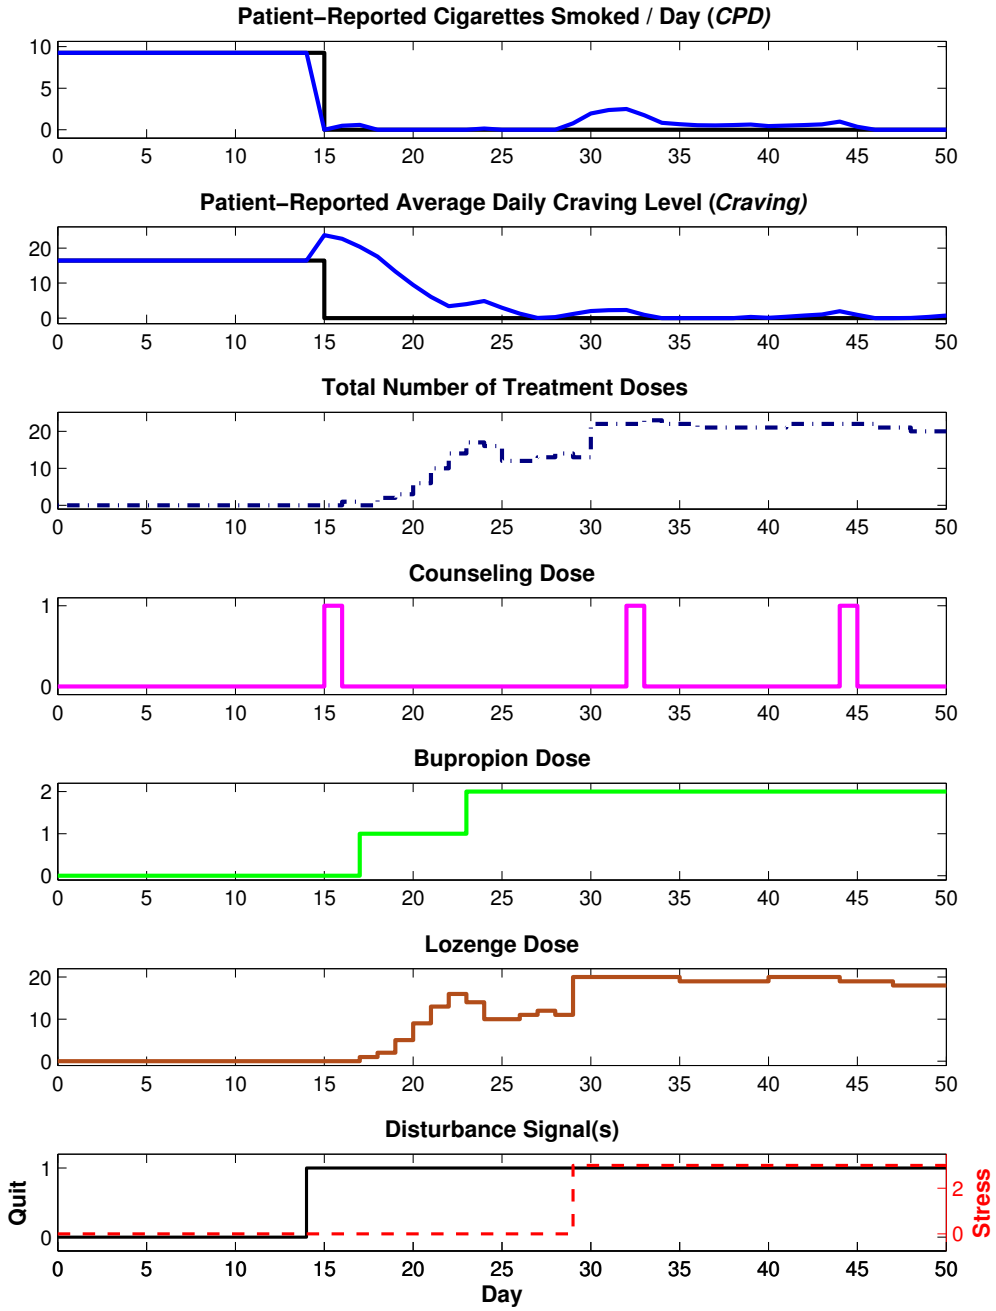

(a) Scenario 3: Outcome and Manipulated Variable Responses in the Presence of a Step Change in the *Stress* Disturbance of Magnitude Three on Day 29 where $Q_{cpd} = 10$, $Q_{crav} = 10$, and $Q_{wip_T} = 0$.

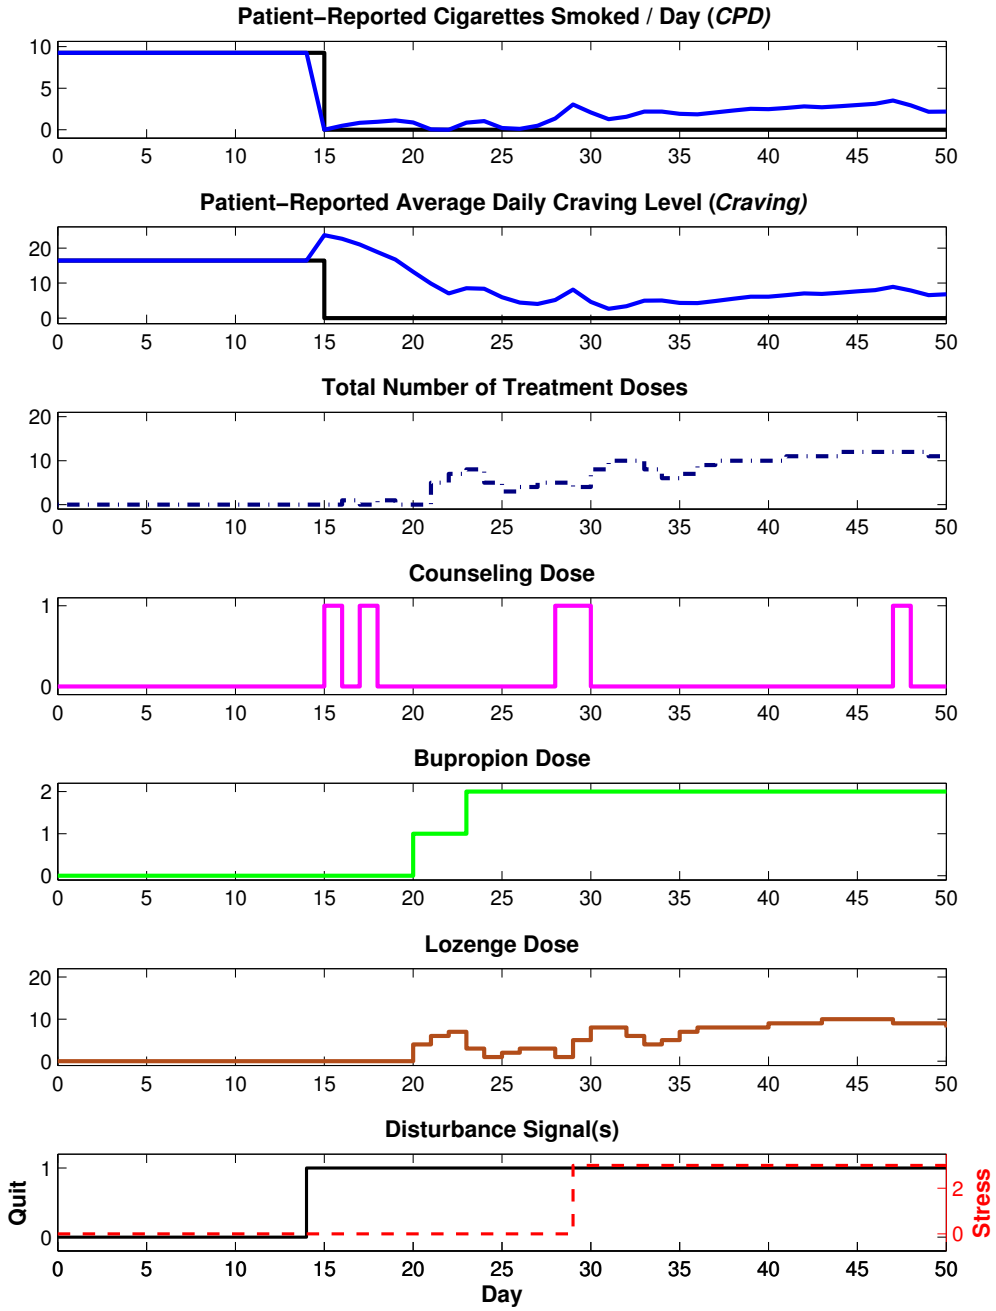

(b) Scenario 3: Outcome and Manipulated Variable Responses in the Presence of a Step Change in the *Stress* Disturbance of Magnitude Three on Day 29 where $Q_{cpd} = 10$, $Q_{crav} = 1$, and $Q_{wip_T} = 1$.

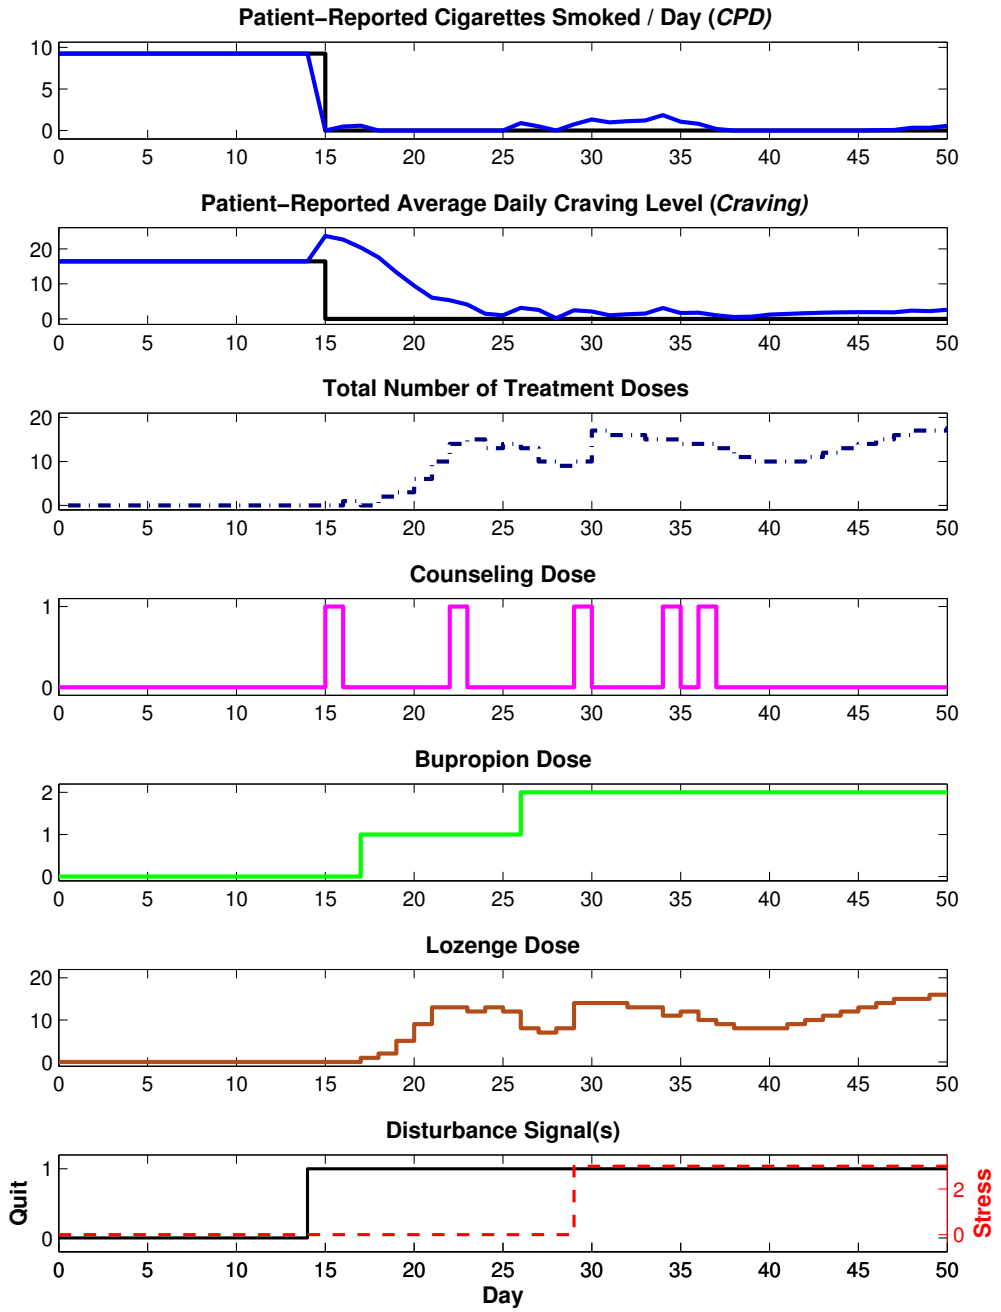

(c) Scenario 3: Outcome and Manipulated Variable Responses in the Presence of a Step Change in the *Stress* Disturbance of Magnitude Three on Day 29 where $Q_{cpd} = 10$, $Q_{crav} = 10$, and $Q_{wip_T} = 1$.

Figure 4.6: Scenario 3: Outcome and Manipulated Variable Responses in the Presence of a Step Change in the *Stress* Disturbance of Magnitude Three on Day 29 where $Q_{cpd} = 10$ for Various Combinations of Q_{crav} and Q_{wip_T} values.

Table 4.4: Scenario 3: Performance Metrics for the Intervention in the Presence of a Step Change in the *Stress* Disturbance of Magnitude Three on Day 29 for Several Combinations of Controlled Variable Penalty Weights.

	Figure:	4.6a	4.6b	4.6c		
Parameters	Q_{cpd}	10	10	10	10	10
	Q_{crav}	1	10	1	10	10
	Q_{wip_T}	0	0	1	1	10
Performance, $t \geq$ TQD	t_r^{cpd}	0	0	0	0	0
	t_r^{crav}	9	12		13	
	e_{50}^{cpd}	0.05	0	2.19	0.57	4.56
	e_{50}^{crav}	0.93	0.76	6.81	2.62	11.31
	Total Cigs	21.56	17.92	61.89	13.25	92.91
	Days CPD=0	12	16	2	17	5
	$max(cpd)$	2.77	2.49	3.52	1.85	4.71
	$max(crav)$	23.7	23.7	23.7	23.7	23.7
	$max(WIP_T)$	23	23	12	18	9
	J_e^{cpd}	26.67	24.81	143.95	13.65	351.51
	J_e^{crav}	2582.44	2189.08	3580.69	2241.94	4309.14
	$J_e^{wip_T}$	11883	11369	2475	5632	1319
	$var(cpd)$	0.39	0.45	1.07	0.25	3.19
	$var(crav)$	55.68	45.75	29.34	41.34	22.88
	$var(WIP_T)$	60.43	55.08	16.65	24.89	11.11
	$J_I^{u_c}$	3	3	5	5	5
	$J_I^{u_b}$	106	118	115	109	85
	$J_I^{u_l}$	10097	9502	1619	4420	804
	$var(u_c)$	0.08	0.08	0.12	0.12	0.12
	$var(u_b)$	0.54	0.32	0.52	0.35	0.68
	$var(u_l)$	44.02	41.77	12.25	18.81	7.42

50% smaller than the nearest comparable J_e^{cpd} values in Table 4.4. This favorable performance is achieved through increased reliance on u_c and early use of u_b. Relatedly, this case features less reliance on u_l dosing and less dramatic changes in daily u_l dosing (see $J_I^{u_l}$, $max(u_l)$, and $var(u_l)$); this less intense lozenge dosing may increase chances of overall adherence (Piper *et al.*, 2009).

Scenario 4

Intuitively, a *Stress* disturbance in the form of a deterministic step signal that remains at a constant level for at least three weeks seems unlikely in real world settings. As noted in Shiffman and Waters (2004), instances of lapse can more directly be attributed to rapidly changing stress levels than to slowly time-varying and chronic stressors. This motivates assessment of controller performance when CPD and *Craving* are subjected to a *Stress* disturbance that is stochastic in nature. For comparative purposes, one realization of an autocorrelated, stochastic stress disturbance is considered initially, the character of which is described by equation 4.1. Fig. 4.7 features intervention responses where $Q_{cpd} = 10$ and varying combinations of Q_{crav} and Q_{wip_T}. The corresponding performance metrics in Table 4.5 offer insight into how varying the penalty weights influences how sensitive control action is to the stochastic stress disturbance, i.e., how aggressive the controller is in terms of rejecting this *Stress* influence, and how that relates to intervention outcomes.

Comparing the performances of the controllers when there is and is not stochastic *Stress* present where $Q_{cpd} = Q_{crav} = 10$ and $Q_{wip_T} = 1$ (see Fig. 4.5b versus 4.7d) indicates that control action, particularly that related to u_l, is sensitive enough to the stochastic disturbance that the patient is still able to achieve relatively successful cessation outcomes. Specifically, there is no CPD offset on day 50 and an additional smoke-free day in the presence of stochastic *Stress* compared to the $Stress = 0$

Table 4.5: Scenario 4: Performance Metrics for the Intervention with Various Penalty Weight Values for a Single Realization of a Stochastic *Stress* Signal (Sto., ex) and Averaged Over 50 Realizations of the Stochastic Signal (Sto., 50).

Stress form		Sto., ex	Sto., ex	Sto., ex	Sto., ex	Sto., ex	Sto., ex	Sto., 50	Sto., 50
	Figure:	4.7a	4.7b	4.7c	4.7d				
Parameters	Q_{cpd}	10	10	10	10	10	10	10	10
	Q_{crav}	0.1	1	10	10	1	10	10	10
	Q_{wip_T}	0	0	0	0.1	1	10	0.1	1
Performance, $t \geq$ TQD	t_r^{cpd}	0	0	0	0	0	0	0.16	0.04
	t_r^{crav}	11	13	11	13	13	20	9.86	13.02
	e_{50}^{cpd}	0	1.13	1.16	1.13	0	0.44	0.24	0.14
	e_{50}^{crav}	1.17	1.35	1.84	1.34	1.8	4.32	2.1	2.97
	Total Cigs	15.03	14.88	23.54	19.79	15.98	40.51	15.41	8.26
	Days CPD=0	19	15	10	13	15	3	14.94	18.22
	$max(cpd)$	2.9	2.9	2.9	2.9	2.9	3.26	1.8	1.79
	$max(crav)$	22.88	22.88	22.88	22.88	22.88	22.88	23.17	22.99
	$max(WIP_T)$	22	21	22	21	17	8	20.24	15.48
	J_e^{cpd}	22.63	18.75	31.87	24.36	24.06	72.23	21.1	8.99
	J_e^{crav}	2564.99	2549.3	2523.96	2521.18	2635.57	3280.32	2207.87	2326.25
	$J_e^{wip_T}$	5525	6577	8226	7742	3358	339	6289.78	2639.6
	$var(cpd)$	0.47	0.36	0.47	0.38	0.48	0.76	0.33	0.19
	$var(crav)$	53.73	54	54.47	54.6	49.9	40.16	45.56	41.7
	$var(WIP_T)$	35.28	39.11	47.57	44.39	22.6	3.4	33.2	16.38
	$J_I^{u_c}$	2	1	2	1	3	5	1.06	2.52
	$J_I^{u_b}$	79	85	58	67	97	58	86.8	103.06
	$J_I^{u_l}$	4489	5497	7127	6710	2417	137	5233.3	1856.34
	$var(u_c)$	0.05	0.03	0.05	0.03	0.08	0.12	0.03	0.07
	$var(u_b)$	0.5	0.52	0.39	0.45	0.54	0.51	0.34	0.43
	$var(u_l)$	29.36	31.97	39.32	36.66	19.06	2.01	27.38	13.41

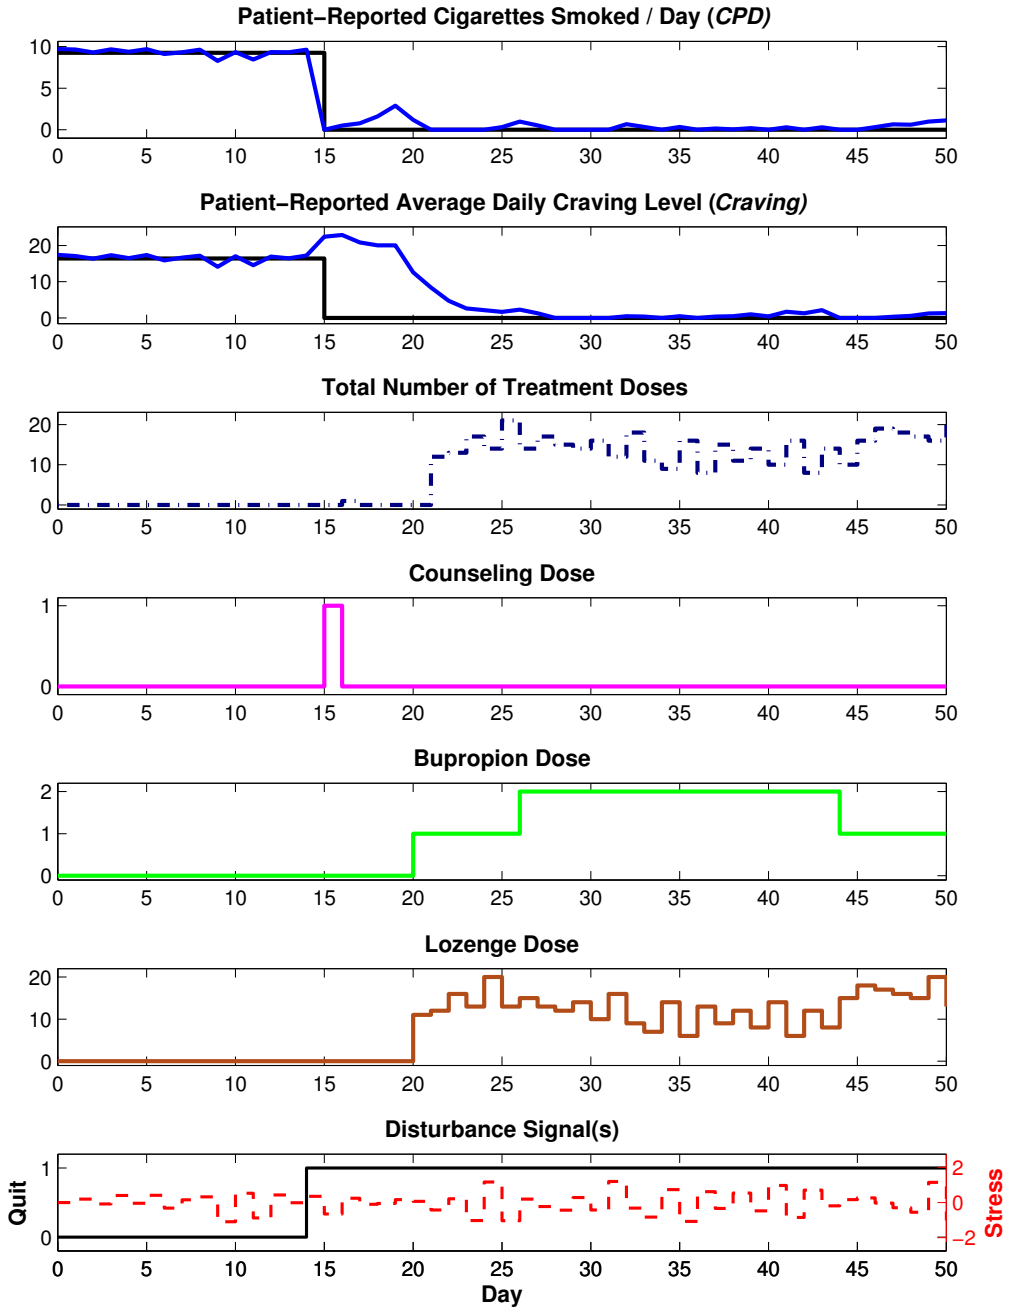

(a) Scenario 4: Outcome and Manipulated Variable Responses where $Q_{cpd} = 10$, $Q_{crav} = 1$, and $Q_{wip_T} = 0$.

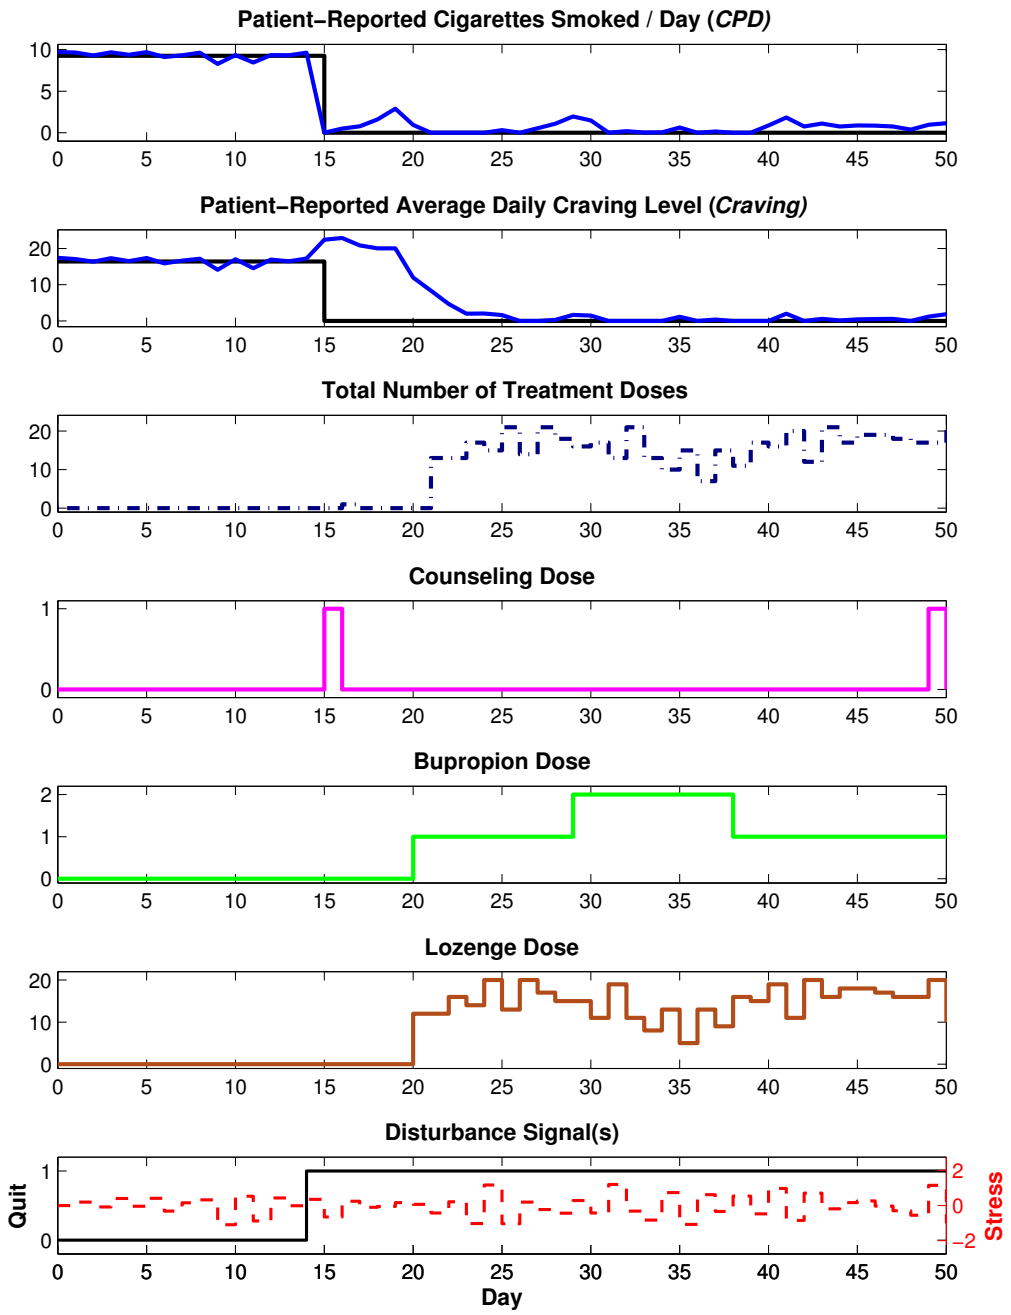

(b) Scenario 4: Outcome and Manipulated Variable Responses where $Q_{cpd} = 10$, $Q_{crav} = 10$, and $Q_{wip_T} = 0$.

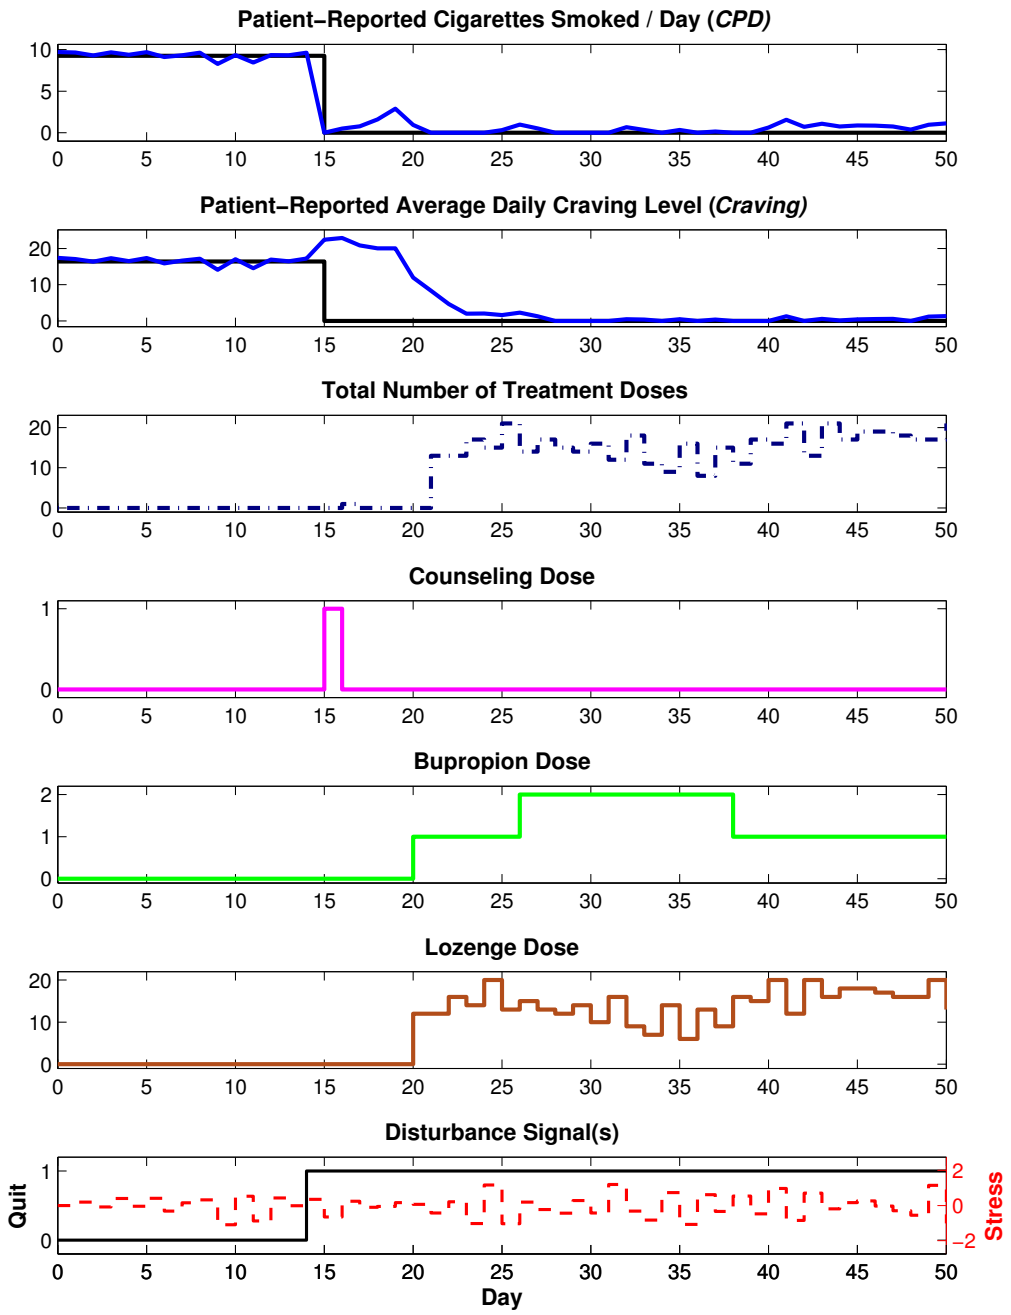

(c) Scenario 4: Outcome and Manipulated Variable Responses where $Q_{cpd} = 10$, $Q_{crav} = 10$, and $Q_{wip_T} = 0.1$.

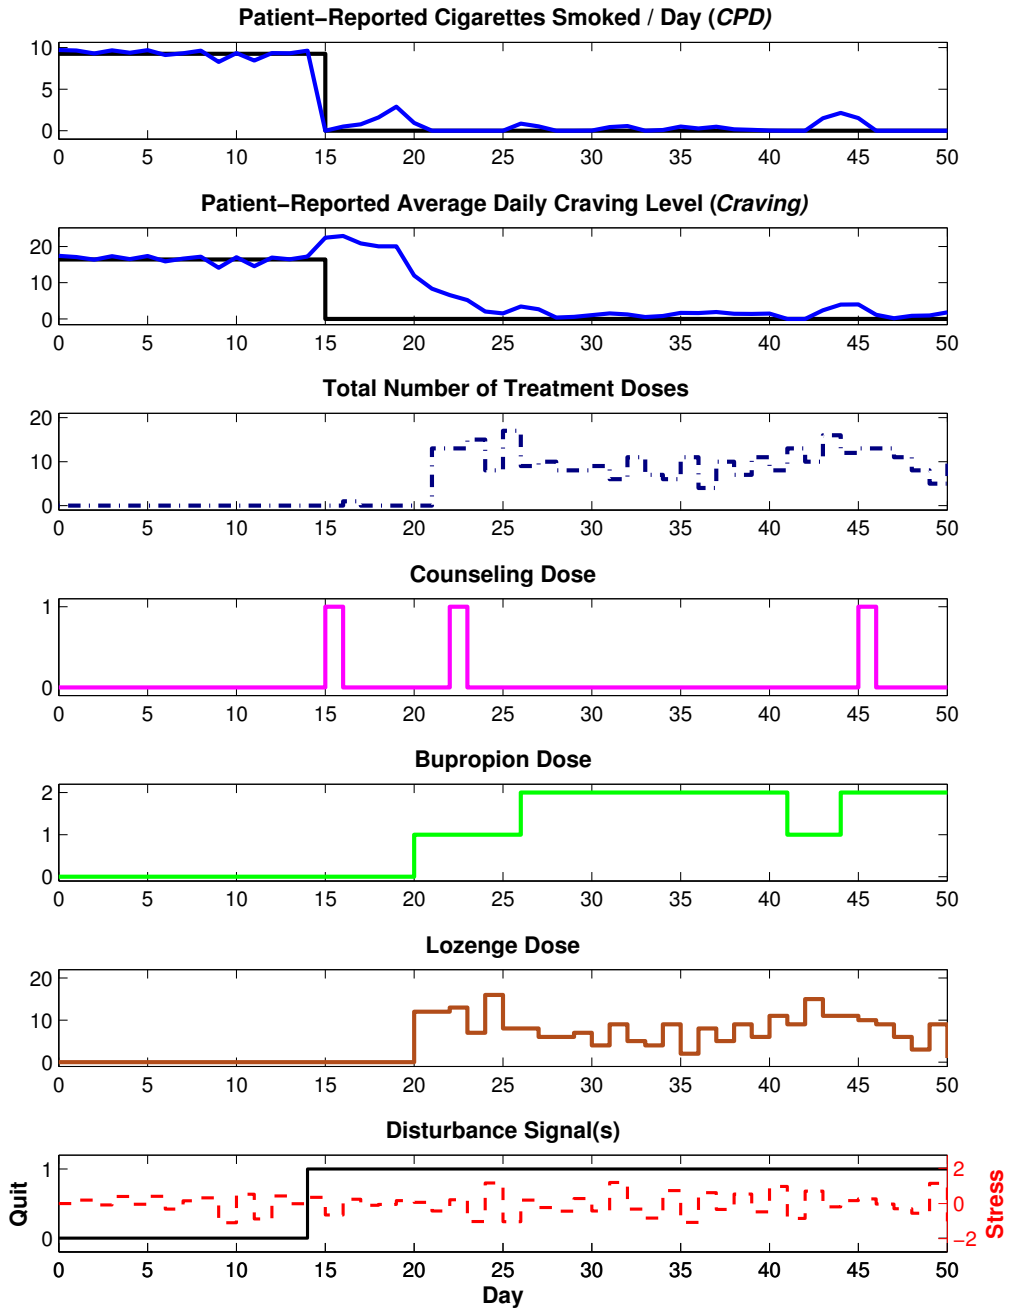

(d) Scenario 4: Outcome and Manipulated Variable Responses where $Q_{cpd} = 10$, $Q_{crav} = 10$, and $Q_{wip_T} = 1$.

Figure 4.7: Scenario 4: Outcome and Manipulated Variable Responses where $Q_{cpd} = 10$ and Q_{crav} and Q_{wip_T} Values Vary. *Stress* is Present in the Form of a Single Stochastic Realization.

case. This performance, though, comes at the expense of dosing requirements as documented in the aggregate dosing metric $J_e^{wip_T}$, which is more than 28 times greater in the stochastic stress case, and dosing variances ($var(WIP_T)$, $var(u_l)$).

Contrasting the scenarios depicted in Fig. 4.6 where $Stress$ is a step disturbance, the controller trades off between u_b and u_l doses in the presence of the stochastic disturbance. As current clinical practice consists of two daily bupropion doses consistently for weeks or months after initially stepping up to two daily doses, decreases in u_b assignments (such as those seen in each scenario in Fig. 4.7) may appear operationally unadvisable to clinicians, regardless of predicted CPD and $Craving$ performances. One potential method for discouraging such dosage changes would be to penalize changes in u_b via the move size penalty $Q_{\Delta u_b}$. Fig. 4.8a depicts the responses and dosing for the case where $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$ and $Q_{\Delta u_b} = 100$. This figure suggests that a large bupropion move size penalty could facilitate an effective intervention that does not require decreases in u_b dosage at any point. As previously described, complex tuning via set point tracking and move suppression weights in the objective function is an unintuitive way to obtain desirable performance characteristics (Nandola and Rivera, 2013). Furthermore, penalizing u_b moves serves to suppress *both* increases and decreases in u_b, which may unnecessarily prevent implementation of efficacious bupropion doses.

Scenario 5

Given the lack of a reliable, systematic way to choose $Q_{\Delta u_b}$ values that penalize $\Delta u_b(k) < 0$ but not $\Delta u_b(k) > 0$, this move size concern is addressed in an alternative manner. Specifically, the move size constraint in equation 3.57 is replaced by equation 4.4.

$$0 \leq \Delta u_b(k) \leq 1 \tag{4.4}$$

Equation 4.4 states that once u_b has increased in dose, it cannot subsequently decrease. Fig. 4.8b depicts the dosing and controlled variable responses with this updated constraint. It indicates adjustment of this constraint does not necessarily degrade performance significantly. For example, $Total\ Cigs$ is only slightly smaller for this $Stress$ realization in the simulation employing the updated formulation, versus the $Q_{\Delta u_b} = 100$ case (although this comparable performance is at the expense of overall dosing requirements, $J_e^{wip_T}$ equal to 3116 versus 2660).

Figures 4.7 and 4.8 examine nominal performance for one realization of the stochastic $Stress$ disturbance. To get a more general picture of the performance subject to $Stress$, Table 4.5 also contains performance metrics for two tuning weight combinations as averaged across the dosing and responses for 50 realizations of the stochastic disturbance. The first set of 50 realizations was obtained for $Q_{cpd} = Q_{crav} = 10$ and $Q_{wip_T} = 0.1$; the second adjusts Q_{wip_T} to equal 1. As one would expect, the realizations with $Q_{wip_T} = 1$ leads to less total dosing on average, represented by an average $J_e^{wip_T}$ that is nearly 2.5 times smaller than that for the $Q_{wip_T} = 0.1$ case. Similarly, the $Q_{wip_T} = 0.1$ case leads to average total dosing variation approximately twice that of the $Q_{wip_T} = 1$ ($var(WIP_T) = 32.7$ versus 16.4). Together, this suggests that increasing the WIP_T objective function penalty weight by an order to magnitude leads to a much less demanding dosing regimen overall. Interestingly, this larger Q_{wip_T} value also corresponds here to improved performances on average, as quantified by lower J_e^{cpd}, $max(cpd)$, $Total\ Cigs$, and $Days\ CPD$=0 metrics. These seemingly incongruous average dosing and CPD performance metrics may be due in part to simulation relaxation. ϵ_T is nearly 50% larger for the $Q_{wip_T} = 1$ case, indicating that more hard constraint relaxation was required for feasible dosing optimizations on average. This may be partially related to more pre-TQD dosing that was assigned more frequently in the $Q_{wip_T} = 1$ case, which would lead to CPD values below 0 on TQD due to the

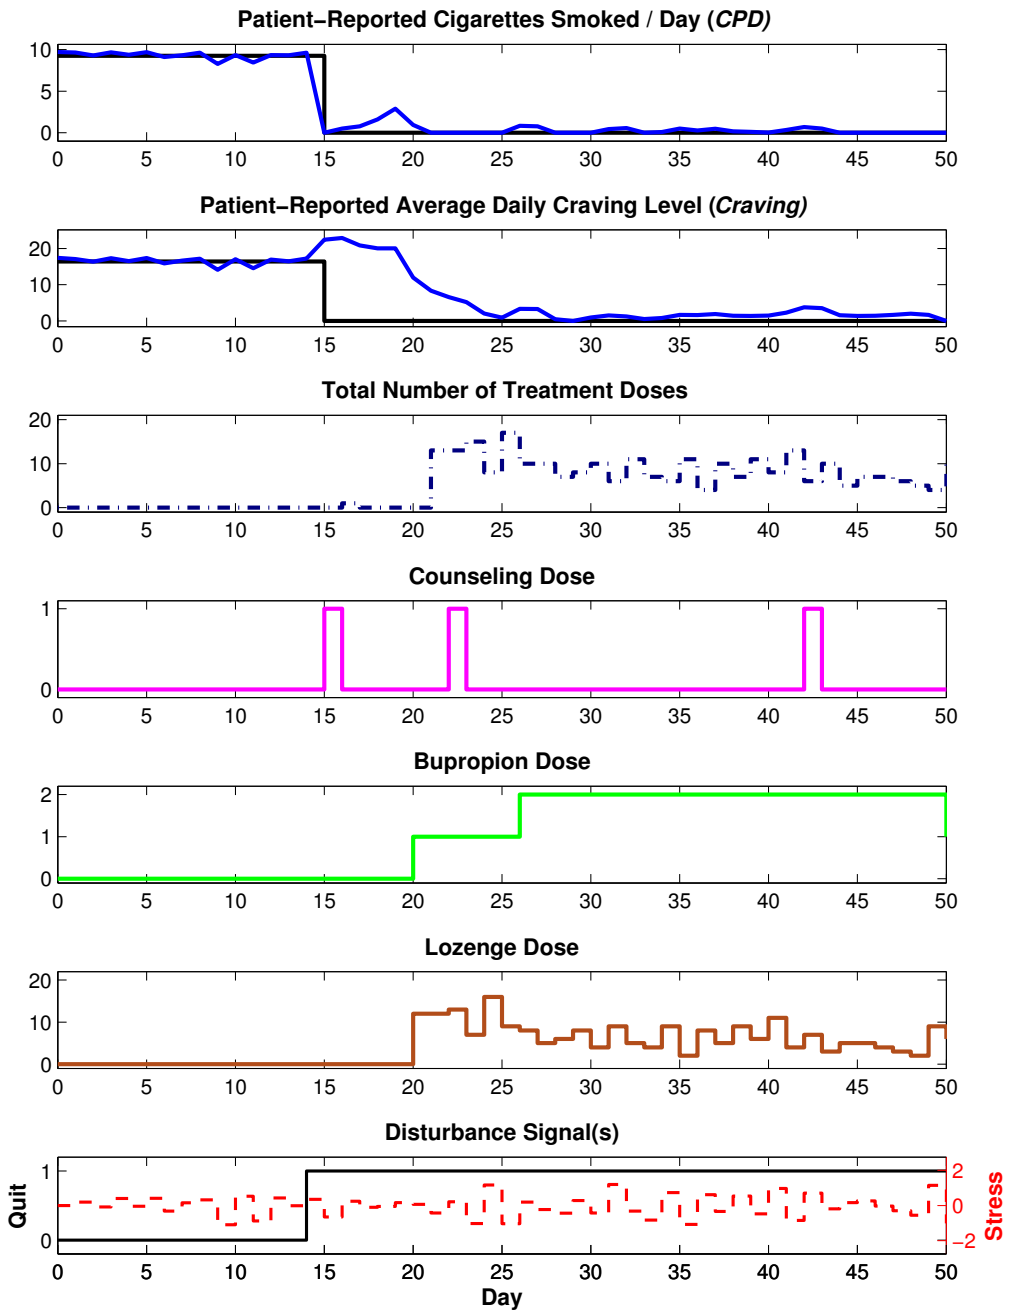

(a) Scenario 4: Outcome and Manipulated Variable Responses where $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$, $Q_{\Delta u_b} > 0$, and the Δu_b Constraint is Described by Equation 3.57.

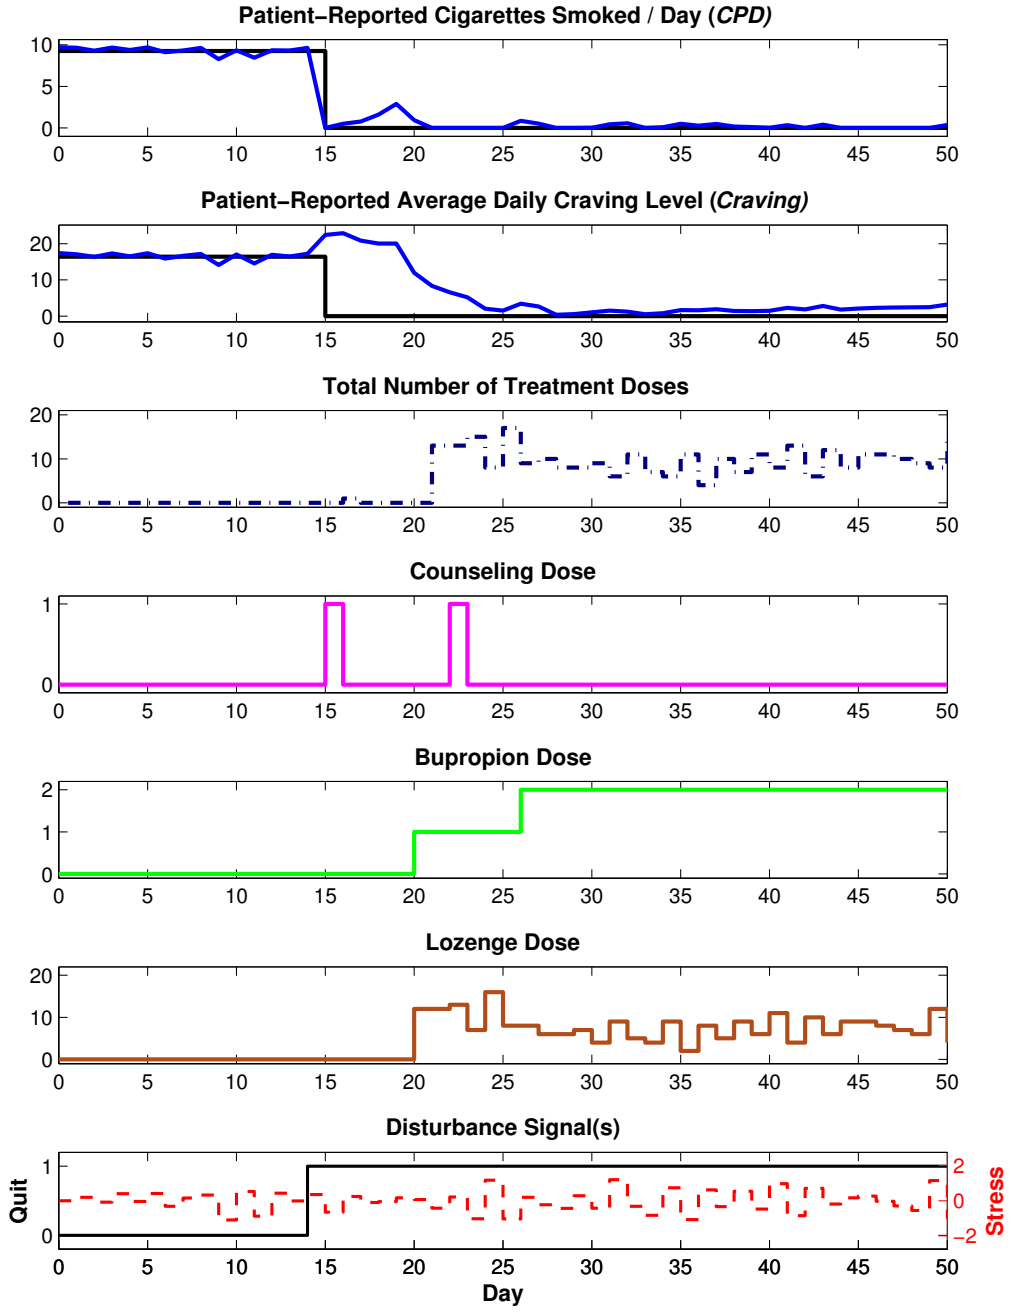

(b) Scenario 5: Outcome and Manipulated Variable Responses where $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$, $Q_{\Delta u_b} = 0$, and the Δu_b Constraint is Described by Equation 4.4.

Figure 4.8: Outcome and Manipulated Variable Responses where $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$, and a Means to Suppress Decreases in u_b Over Time.

Quit disturbance. Pre-TQD dosing indicates that the controller takes action to keep CPD and $Craving$ levels at average baseline values, essentially avoiding conditions immediately prior to TQD that make abstinence early in the quit attempt more difficult; as more of this pre-TQD dosing occurs in with $Q_{wip_T} = 1$, the controller appears to take this action prior to TQD in order to require less dosing at future points in the prediction horizon.

4.4.2 Tuning via 3DoF Tuning Parameters

The focus of Section 4.4.2 remains to be nominal performance, but in the context of the intervention architecture depicted in Fig. 3.5. Under nominal conditions, adjustment of f_a^{cpd} and f_a^{crav} offer little meaningful effect on the manipulated variable assignments or controlled variables responses. Therefore, $f_a^{cpd} = f_a^{crav} = 1$ in the simulations immediately following; detuning via f_a tuning parameters will not be described until robustness is considered in subsequent sections.

Scenario 6

Section 4.4.1 illustrated that without 3DoF functionality, more than marginal alterations in performance—as quantified by the summed total of CPD and $Craving$ post-TQD—require orders of magnitude of Q_{cpd} and Q_{crav} adjustment. This finding suggests that tuning via α_r^{cpd} or α_r^{crav} independently would need to be aggressive in order to to nontrivially affect the controlled variable offset magnitudes, as summed across the quit attempt time frame considered here (for a given set of penalty weights and in the absence of additional tuning parameters). Fig. 4.9 depicts relative J_e^{cpd}, J_e^{crav}, and $J_e^{wip_T}$ values as a function of α_r^{cpd} and α_r^{crav} and confirms these expectations.

Examining the α_r^{cpd} and α_r^{crav} axes in Fig. 4.9a indicates that even large, independent adjustments of either α_r^* parameter does not lead to large changes in total

post-TQD smoking. However, this figure reflects that simultaneous adjustment of the α_r^* parameters such that $\alpha_r^{cpd} > 0.5$ and $\alpha_r^{crav} > 0.6$ leads to much greater magnitudes of smoking during the first 5 weeks of a quit attempt. For example, when $\alpha_r^{cpd} = \alpha_r^{crav} = 0.8$, the patient smokes more than twice as many cigarettes than the no-3DoF tuning case over the same time period (*Total Cigs* in Table 4.6), with the peak lapse event having a magnitude more than double that of the no-tuning case ($max(cpd)$). Negligible changes in the summed total of post-TQD $Craving$ are observed for cases where $\alpha_r^{crav} \leq 0.6$. However, Fig. 4.9b indicates that even as α_r^{crav} increases beyond 0.6, total post-TQD $Craving$ increases marginally compared to the scenario with very aggressive tracking of the $Craving$ target. As would be expected, Fig. 4.9c, the total amount of post-TQD dosing reflects aggressive pursuit of the CPD and $Craving$ targets in the absence of significant detuning. Specifically, the values of $J_e^{wip_T}$ are marginally affected by $\alpha_r^{crav} \leq 0.6$, although even more extreme detuning of the $Craving$ target tracking has little influence on the peak level of total daily dosing.

Altogether, Fig. 4.9 indicates that detuning the intervention targets for the purposes of lowering total dosing requirements should focus on $\alpha_r^{cpd} < 0.6$ and $\alpha_r^{crav} > 0.6$, and that the engineer or practitioner should be most concerned with the CPD outcome when altering these parameters. Examining this α_r^* window, the average dosing metrics in Table 4.6 suggest an $\alpha_r^{crav} \approx 0.8$ may correspond to more appealing dosing. Specifically, the $\alpha_r^{crav} = 0.8$ and $\alpha_r^{cpd} = 0$, 0.4, or 0.6 cases require additional bupropion doses, 20 to 33 fewer lozenges, and less intense lozenge adjustments ($var(u_l)$) compared to the other tuning combinations. These tuning cases also correspond to an additional cigarette-free day and a 24% smaller peak lapse event.

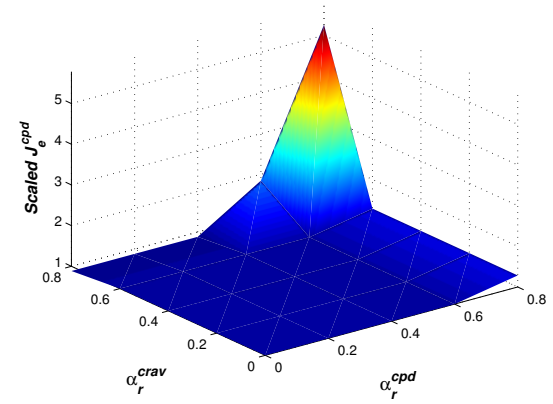

(a) Relative J_e^{cpd} Values.

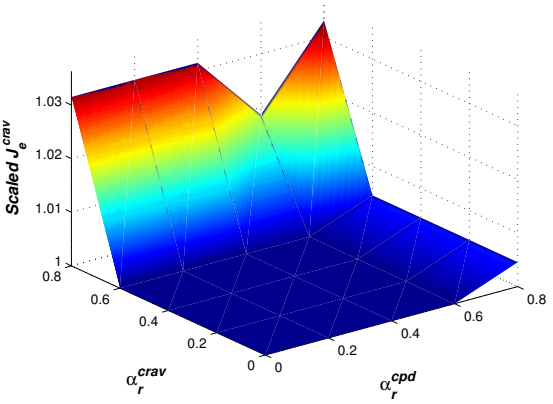

(b) Relative J_e^{crav} Values.

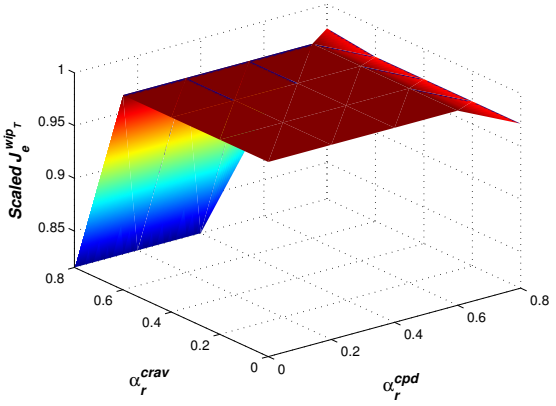

(c) Relative $J_e^{wip_T}$ Values.

Figure 4.9: Scenario 6: Nominal, $Stress = 0$, Post-TQD Intervention Performance Metrics for $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$ and Varied α_r^{cpd} and α_r^{crav} Parameters Relative to the $\alpha_r^{cpd} = \alpha_r^{crav} = 0$ Case.

Table 4.6: Scenario 6: Performance Metrics for the Intervention When $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$ for Various Levels of Detuning via α_r^{cpd} and α_r^{crav} ($Stress = 0$).

Parameters	α_r^{cpd}	0.2	0.4	0.8	0	0	0	0.2	0.4	0.6	0.8
	α_r^{crav}	0	0	0	0.2	0.4	0.8	0.2	0.4	0.6	0.8
Performance, $t \geq$ TQD	t_r^{cpd}	0	0	0	0	0	0	0	0	0	0
	t_r^{crav}	13	13	13	13	13		13	13	13	16
	e_{50}^{cpd}	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06
	e_{50}^{crav}	2.72	2.72	2.72	2.72	2.72	2.72	2.72	2.72	2.72	2.72
	Total Cigs	5.55	5.55	5.9	5.55	5.55	5.45	5.55	5.55	5.55	11.42
	Days CPD=0	14	14	14	14	14	15	14	14	14	10
	$max(cpd)$	0.91	0.91	1.19	0.91	0.91	0.69	0.91	0.91	0.91	2.1
	$max(crav)$	23.7	23.7	23.7	23.7	23.7	23.7	23.7	23.7	23.7	23.7
	$max(WIP_T)$	15	15	15	15	15	15	15	15	15	14
	J_e^{cpd}	2.26	2.26	2.9	2.26	2.26	2	2.26	2.26	2.26	11.89
	J_e^{crav}	2248	2248	2258	2248	2248	2319	2248	2248	2248	2322
	$J_e^{wip_T}$	3080	3080	2992	3080	3080	2512	3080	3080	3080	3010
	$var(cpd)$	0.04	0.04	0.06	0.04	0.04	0.03	0.04	0.04	0.04	0.24
	$var(crav)$	41.48	41.48	41.01	41.48	41.48	40.43	41.48	41.48	41.48	40.42
	$var(WIP_T)$	12.71	12.71	12.14	12.71	12.71	10.43	12.71	12.71	12.71	11.69
	$J_I^{u_c}$	2	2	2	2	2	2	2	2	2	2
	$J_I^{u_b}$	109	109	109	109	109	118	109	109	109	100
	$J_I^{u_l}$	2173	2173	2097	2173	2173	1654	2173	2173	2173	2150
	$var(u_c)$	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05
	$var(u_b)$	0.35	0.35	0.35	0.35	0.35	0.32	0.35	0.35	0.35	0.37
	$var(u_l)$	9.67	9.67	9.11	9.67	9.67	7.43	9.67	9.67	9.67	8.6

Scenario 7

In the presence of a measured disturbance, tuning should first adjust α_d parameters prior to α_r^* parameter adjustment, as the manner in which controller targets are obtained and maintained is a function of the disturbance rejection capabilities of the controller. As before, the *Stress* disturbance takes in the form of a step of magnitude three two weeks into the quit attempt, and in the form of a stochastic signal beginning on day 1. Initially, independent tuning of the α_d^S parameter is examined.

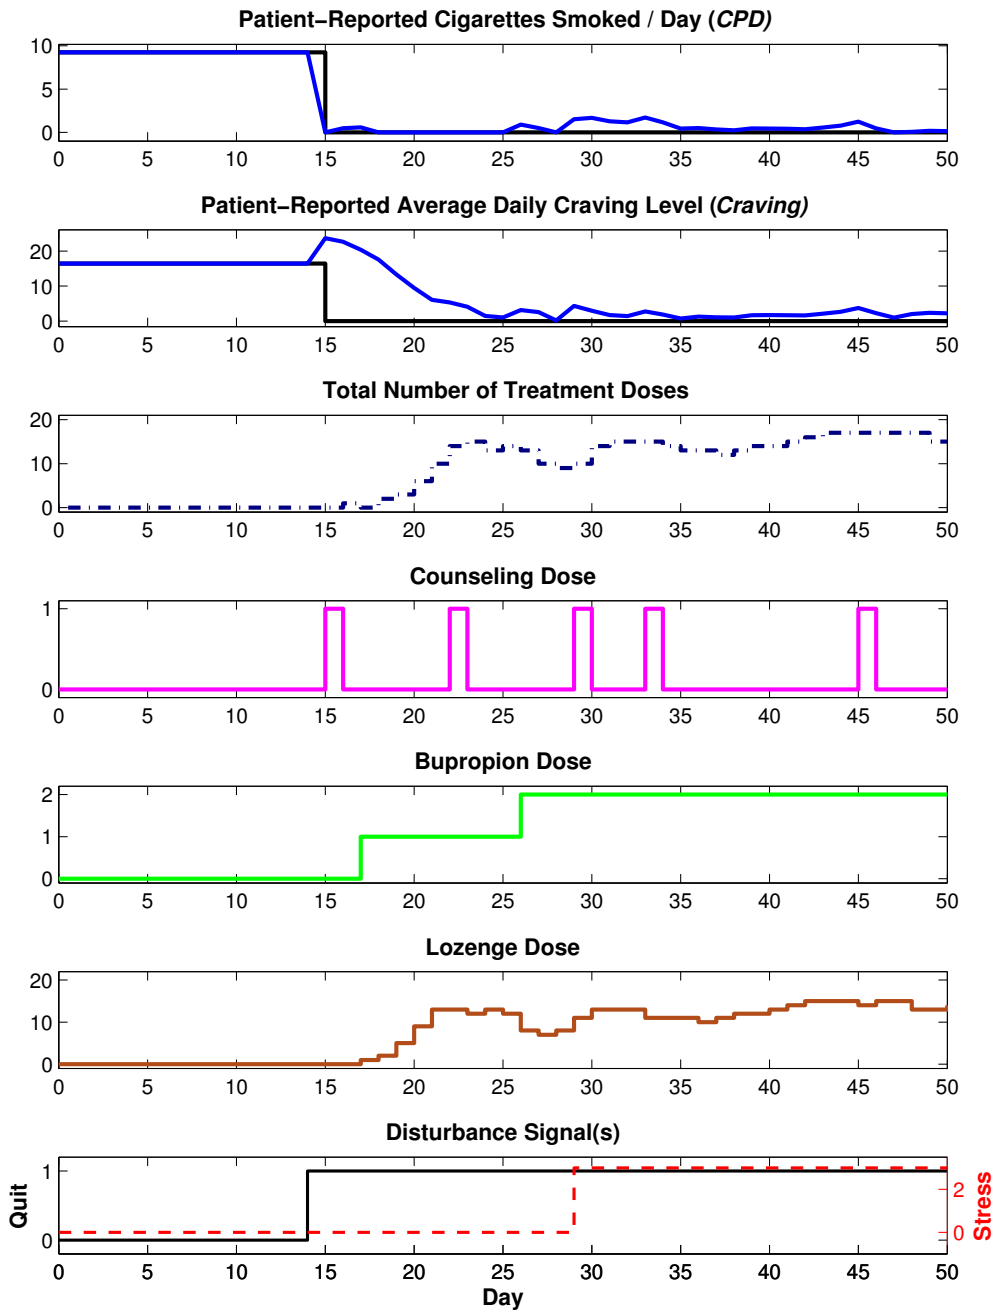

(a) Scenario 7: Outcome and Manipulated Variable Responses where $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$, $\alpha_r^{cpd} = \alpha_r^{crav} = 0$, and $\alpha_d^S = 0.2$.

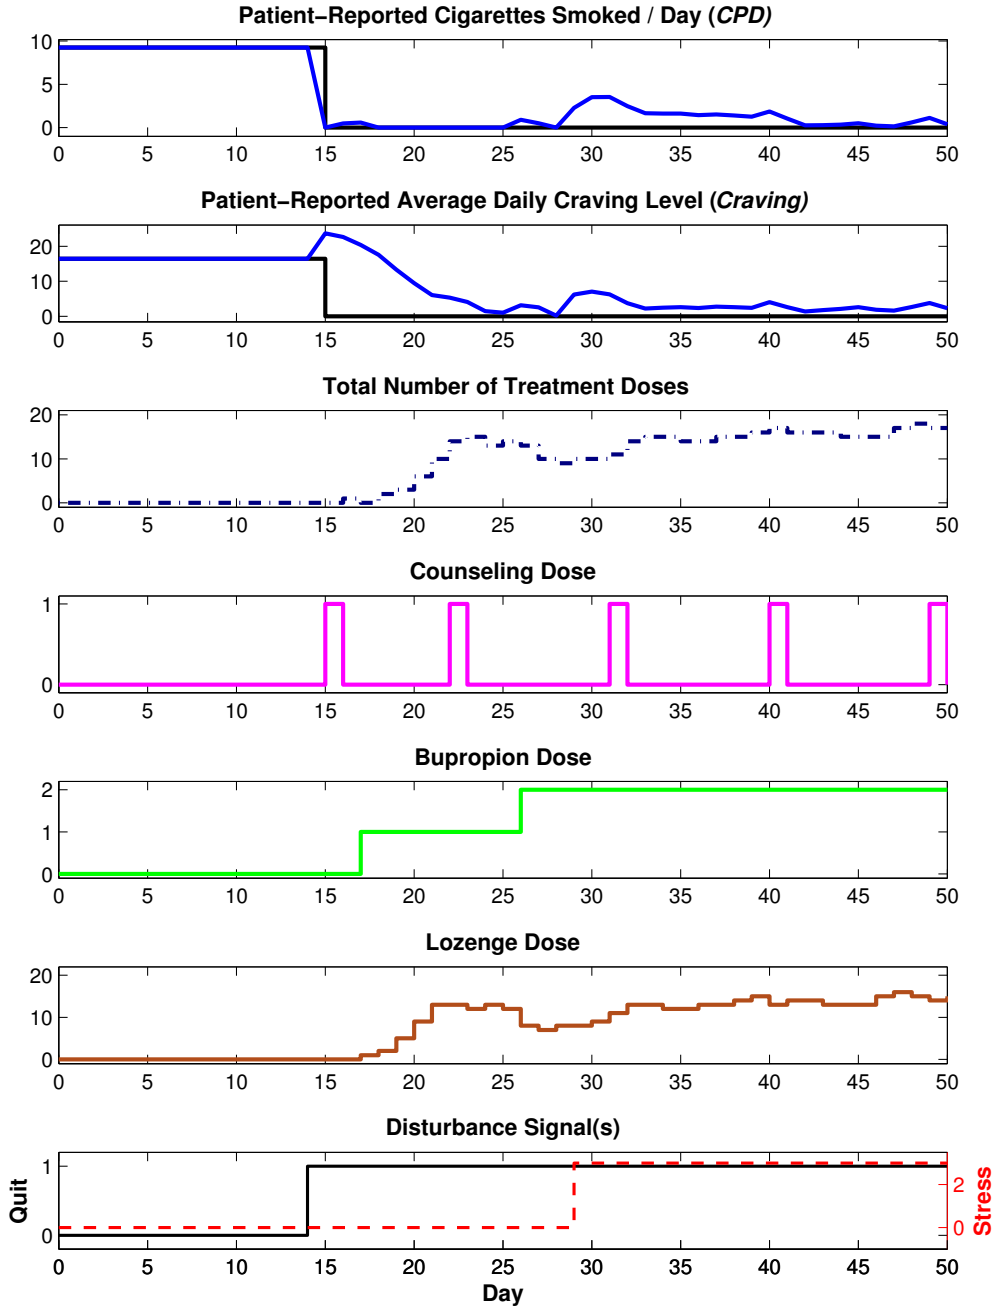

(b) Scenario 7: Outcome and Manipulated Variable Responses where $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$, $\alpha_r^{cpd} = \alpha_r^{crav} = 0$, and $\alpha_d^S = 0.8$.

Figure 4.10: Scenario 7: Outcome and Manipulated Variable Responses where $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$, $\alpha_r^{cpd} = \alpha_r^{crav} = 0$, and Various α_d^S Values.

Table 4.7: Scenario 7: Performance Metrics for the Intervention Under a Selection of 3DoF Tuning Conditions Where $\alpha_d^S > 0$ and *Stress* Takes the Form of a Step of Magnitude Three on Day 29.

	Figure:		4.10a	4.10a
Parameters	α_r^{cpd}	0	0	0
	α_r^{crav}	0	0	0
	α_d^S	0	0.2	0.8
Performance, $t \geq$ TQD	t_r^{cpd}	0	0	0
	t_r^{crav}	13	13	13
	e_{50}^{cpd}	0.21	0.15	0.38
	e_{50}^{crav}	2.31	2.21	2.31
	Total Cigs	15.46	17.77	31.64
	Days CPD=0	12	11	10
	$max(cpd)$	1.65	1.73	3.54
	$max(crav)$	23.7	23.7	23.7
	$max(WIP_T)$	18	17	18
	J_e^{cpd}	13.2	18.29	60.84
	J_e^{crav}	2249.22	2273.38	2426.15
	$J_e^{wip_T}$	6367	6161	6274
	$var(cpd)$	0.19	0.27	0.94
	$var(crav)$	41.29	40.55	38
	$var(WIP_T)$	27.56	25.85	27
	$J_I^{u_c}$	5	5	5
	$J_I^{u_b}$	109	109	109
	$J_I^{u_l}$	4979	4803	4929
	$var(u_c)$	0.12	0.12	0.12
	$var(u_b)$	0.35	0.35	0.35
	$var(u_l)$	19.68	18.36	19.49

Table 4.7 and Fig. 4.10 suggest that for a measured and anticipated step disturbance of this character, disturbance rejection capabilities are most significantly degraded with large α_d^S values (in the absence of additional detuning). That said, the CPD and $Craving$ offsets 21 days after the step disturbance are only 0.38 cigarettes and 2.31 points, respectively, when $\alpha_d^S = 0.8$. This CPD offset is only 77% larger than that for the $\alpha_d^S = 0$ case and less than 0.4 cigarettes greater than the lowest e_{50}^{cpd} value in Table 4.7, while this $Craving$ offset is approximately equal to that for the $\alpha_d^S = 0$ case and less than 1 point larger than the smallest e_{50}^{crav} magnitude. However, significantly slowing rejection of the step disturbance leads to fewer cigarette-free days (2 fewer days for $\alpha_d^S = 0.8$).

Overall, though, post-TQD CPD and $Craving$ responses are negatively affected as α_d^S increases; *Total Cigs*, J_e^{cpd}, J_e^{crav}, $max(cpd)$, var(cpd), and var(crav) values in Table 4.7 are all more than double for the large detuning case. Interestingly, such degraded performance occurs despite relatively similar total dosing demands, as reflected in $J_e^{wip_T}$ and manipulated variable variance metrics. The $\alpha_d^S = 0.4$ case appears to be more responsive to the disturbance initially, assigning two u_c doses within five days of the step, and relatively fast changes in u_l assignment over a similar time period. Altogether, these simulations suggest that detuning via α_d^S may not offer an acceptable trade off between dosing and performance degradation in the presence of a step disturbance of this nature.

Scenario 8

Fig. 4.11 depicts the average cumulative energy metrics J_e^{cpd}, J_e^{crav}, and $J_e^{wip_T}$ from 25 realizations of the stochastic disturbance signal (scaled to the average J_e^* values for the $\alpha_d^S = 0$ case) where the intervention is detuned by increasing values of α_d^S.

When *Stress* takes the form of the stochastic signal described in Section 4.2.2,

Table 4.8: Scenario 8: Performance Metrics for the Intervention Under a Selection of 3DoF Tuning Conditions Where *Stress* Takes the Form of a Stochastic Signal. The Reported Metrics are Averaged Across 25 Realizations of the Disturbance.

Parameters	α_r^{cpd}	0	0	0	0.65	0.8	0	0	0.65
	α_r^{crav}	0	0	0	0	0	0.65	0.8	0.65
	α_d^S	0	0.4	0.8	0.2	0.2	0.2	0.2	0.2
Performance, $t \geq$ TQD	t_r^{cpd}	0.2	0	0.1	0.2	0	0.2	0.1	0.3
	t_r^{crav}	12.2	12	12.8	13.35	13.3	13.1	14.3	13.4
	e_{50}^{cpd}	0.11	0.09	0.15	0.13	0.04	0.14	0.09	0.16
	e_{50}^{crav}	2.88	2.67	2.58	2.8	2.67	2.95	2.84	3.18
	Total Cigs	8.3	8.87	8.54	7.85	9.23	6.11	8.16	7.12
	Days CPD=0	18.84	18.6	18.48	16.55	16.3	19.7	14.5	18.4
	$max(cpd)$	1.79	1.75	1.52	1.69	1.65	1.4	1.65	1.22
	$max(crav)$	23.26	23.29	23.88	23.96	23.6	23.42	23.83	24.1
	$max(WIP_T)$	15.36	15.36	15.36	16.15	16.1	15.4	13.6	14.3
	J_e^{cpd}	9.25	9.25	7.73	7.8	9.17	6.44	8.04	5.97
	J_e^{crav}	2334.79	2363	2299.95	2356.53	2378.65	2243.01	2472.43	2311.92
	$J_e^{wip_T}$	2547.36	2783	2941.12	2823	2831.2	2399.2	2288.2	2472.5
	$var(cpd)$	0.2	0.19	0.16	0.17	0.18	0.14	0.17	0.12
	$var(crav)$	41.77	42.7	41.83	42.06	42.64	39.95	42.48	40.55
	$var(WIP_T)$	16.49	14.46	14.25	16.44	15.42	13.84	12.62	13.75
	$J_I^{u_c}$	2.6	2.32	2.24	2.3	2.3	2.4	2.2	2.6
	$J_I^{u_b}$	102.88	105.04	106.84	108.25	102.7	110.2	110.2	104.8
	$J_I^{u_l}$	1788.6	1949.92	2074.44	1999.35	2023.7	1622.8	1518.4	1704.4
	$var(u_c)$	0.07	0.06	0.06	0.06	0.06	0.06	0.06	0.07
	$var(u_b)$	0.41	0.41	0.36	0.39	0.4	0.38	0.43	0.38
	$var(u_l)$	13.92	11.17	11.17	13.67	12.43	11.24	9.43	11.29

Fig. 4.11 and Table 4.8 suggests that the CPD intervention target is much more sensitive to detuning via α_d^S than is the $Craving$ target. This is consistent with earlier findings suggesting that overall dosing decisions are heavily biased toward the CPD target, which is equivalently the lower bound on CPD. Specifically, relatively minor adjustments to α_d^S or large adjustments to α_d^S—i.e., $\alpha_d^S \leq 0.2,\ > 0.5$—actually improves performance in terms of total smoking during the quit attempt compared

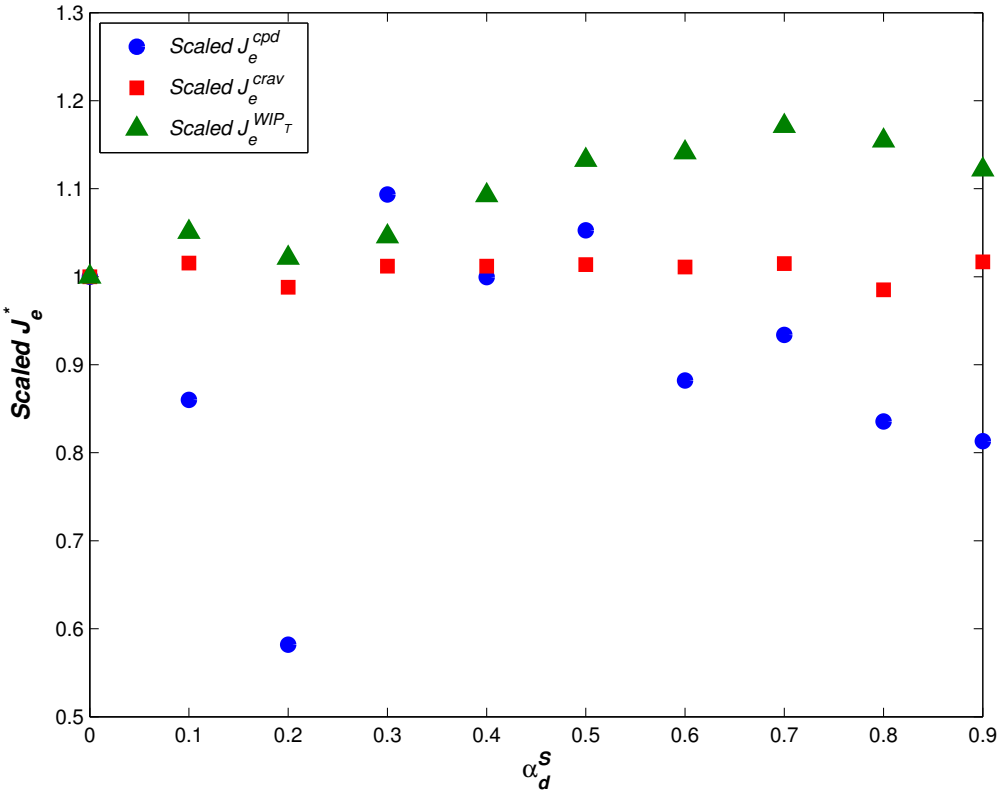

Figure 4.11: Scenario 8: Post-TQD Intervention Performance Metrics (as Averaged Over 25 Stochastic Disturbance Realizations) as Quantified by Total Deviations from Set Points (J_e^*, where $*$ Indicates CPD, $Craving$, or WIP_T) in the Presence of Stochastic Stress for Increasing α_d^S Values, Scaled Relative to J_e^* for the $\alpha_d^S = 0$ Case.

to the most aggressive case. Meanwhile, these α_d^S ranges have little effect on the summed total of post-TQD $Craving$.

Improvement in smoking abstinence success with relatively minor detuning ($0 < \alpha_d^S \leq 0.2$) may result from the autocorrelated nature of the stochastic disturbance signal. As dosing decisions are based in a measured disturbance's forecasted effect, a first order filter on the measured disturbance means dosing decisions are based in a filtered representation of $Stress$. According to equation 3.94, this representation of $Stress$ is a function of both the most recent measured value of $Stress$ and the filtered representation of $Stress$ at the previous time point, $Stress_{flt}(k-1)$. This suggests

that for $\alpha_d^S > 0$, dosing decisions account for a relationship between $Stress(k)$ and $Stress(k-1)$. As the stochastic $Stress$ disturbance considered here features autocorrelation with a first order lag (see equation 4.1), the filter essentially leads to dosing decisions that implicitly recognize the autocorrelated nature of $Stress$, and can therefore anticipate and dose more effectively. As equation 4.1 indicates relatively modest autocorrelation of the stochastic disturbance signal specifically, more effective tracking of the CPD target during the quit attempt overall is understandable for a small window of minor detuning, such as that for $0 < \alpha_d^S \leq 0.2$ (see Fig. 4.11).

Increased total dosing is likely responsible for improvements in total post-TQD smoking levels observed for $\alpha_d^S > 0.5$ (see the scaled J_e^{cpd} and $J_e^{wip_T}$ averages in Fig. 4.11). Interestingly, Fig. 4.11 suggests that the tuning parameter window $0.2 < \alpha_d^S \leq 0.5$ corresponds to no improvement of, or even degraded, performance in terms of total smoking during the quit attempt. This may define a tuning window where the filter serves to suppress increases in dosing that would otherwise occur, but at the expense of set point tracking.

Examining the average dosing demands specifically, decreasing the speed of $Stress$ rejection appears to mitigate the intensity of daily dosing adjustments, as expected. This is reflected by the averaged $var(J_e^{wip_T})$, $var(u_b)$, and $var(u_l)$ values in Table 4.8, which trend downward with increasing α_d^S values. Also suggested by previous tuning scenarios, particularly favorable tuning scenarios may be those that elicit increased use of u_b. Per Table 4.8, $\alpha_d^S = 0.2$ employs the greatest number of bupropion doses over the quit attempt on average, 59, assigning u_b levels more consistently than the other tuning scenarios with comparable total dosing. On average, this tuning scenario and resulting average bupropion dosing led to the smallest peak instance of lapse and least sporadic lapse trends versus the comparable scenarios in Table 4.8, as well as large numbers of smoke-free days on average.

Scenario 9

Given the promising effects of $\alpha_d^S = 0.2$, nominal performance with stochastic $Stress$ is examined where α_d^S fixed to 0.2, and varied levels of filtering of the CPD and $Craving$ targets: recalling that previous adjustment of α_r^{cpd} and α_r^{crav} suggested that $\alpha_r^{cpd} < 0.6$ and $\alpha_r^{crav} > 0.6$ lead to reduced total dosing demands (see Fig. 4.9) and that α_r^{cpd} and α_r^{crav} values below 0.8 have relatively minor effects on post-TQD CPD and $Craving$ performances (see Fig. 4.9a and 4.9b), Table 4.8 documents nominal performance metrics for relatively large α_r^{crav} values. These metrics consist of values averaged across 50 realizations of the stochastic $Stress$ signal. Fig. 4.12 features the associated responses and dosing demands for individual $Stress$ realizations.

As expected, relatively extreme detuning of tracking of the $Craving$ target, $\alpha_r^{crav} = 0.8$ corresponds to the least aggressive dosing documented in Table 4.8. This is reflected in the relatively low values of the u_b variance, u_l variance, $J_e^{wip_T}$, and $J_I^{u_l}$ values. This, though, comes at the expense of performance, highlighted by this scenario's five fewer cigarette-free days on average—likely indicating an undesirable trade-off. Considering the $\alpha_r^{cpd} = \alpha_r^{crav} = 0.65$ and $\alpha_r^{cpd} = 0$, $\alpha_r^{crav} = 0.65$ cases, the former requires smaller peak daily dosing ($max(wip_T) = 14.3$ versus 15.4) and peak lapse levels (max(cpd) = 1.22 cigarettes versus 1.40) on average; however, the latter features approximately 1.5 additional cigarette-free days while requiring fewer total doses on average. In such cases where the performance and dosing trade off is relatively subtle, the appropriateness of one tuning strategy over another may be clarified through a clinician's assessment of a patient's likelihood to adhere to high daily dosing demands or total daily dosing commands.

Altogether, Table 4.8 suggests that detuning via 3DoF functionality where $\alpha_r^{cpd} = 0$, $\alpha_r^{crav} = 0.65$, and $\alpha_d^S = 0.2$ corresponds to favorable intervention performance in

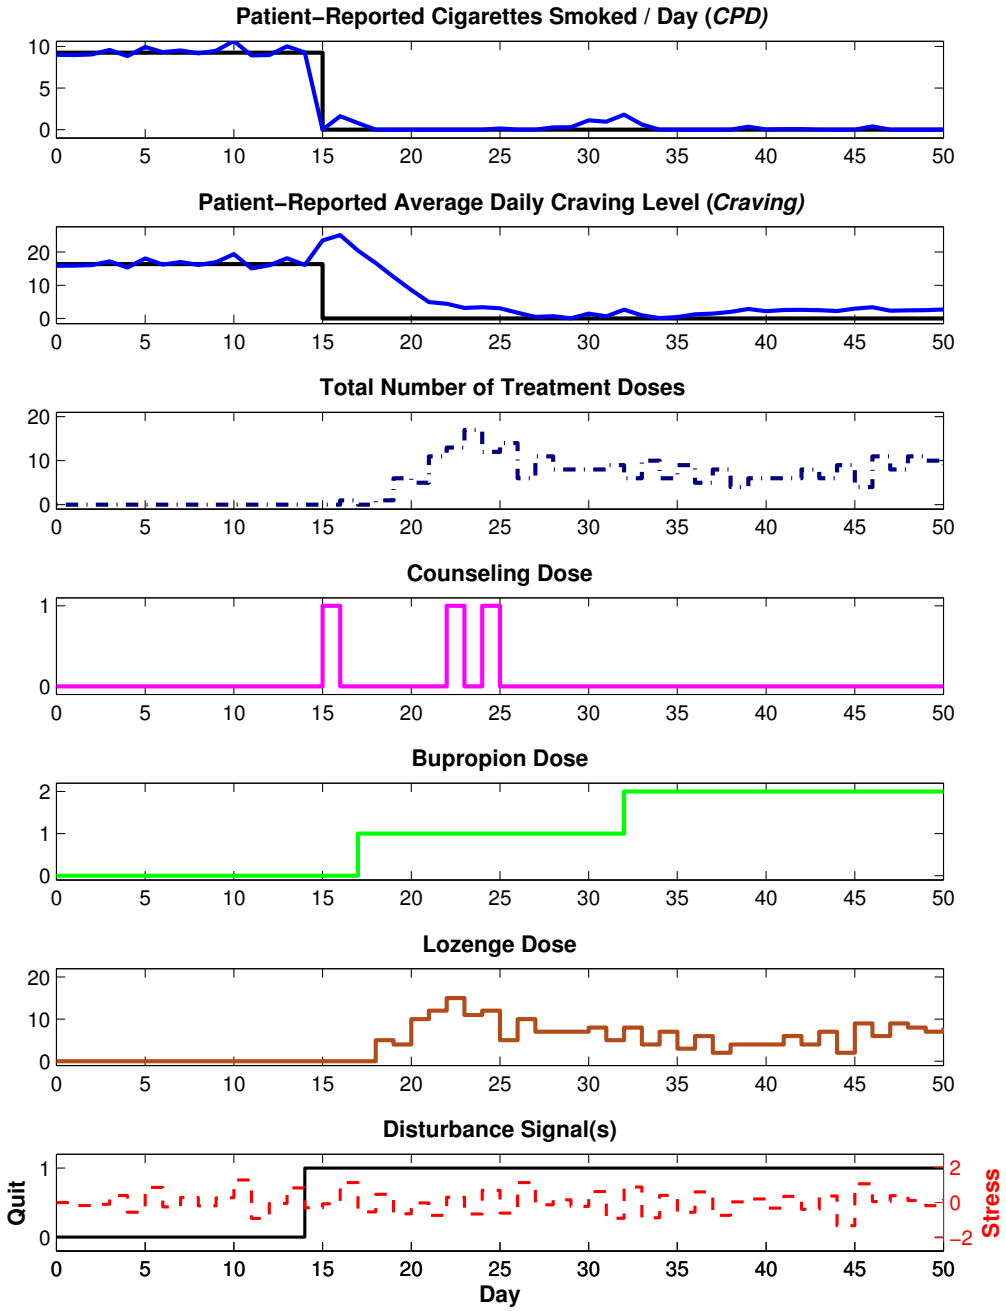

(a) Scenario 9: Outcome and ManipulatedVariable Responses when a Stochastic *Stress* Disturbance is Present and $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$, $\alpha_d^S = 0.2$, $\alpha_r^{cpd} = 0$, and $\alpha_r^{crav} = 0.65$.

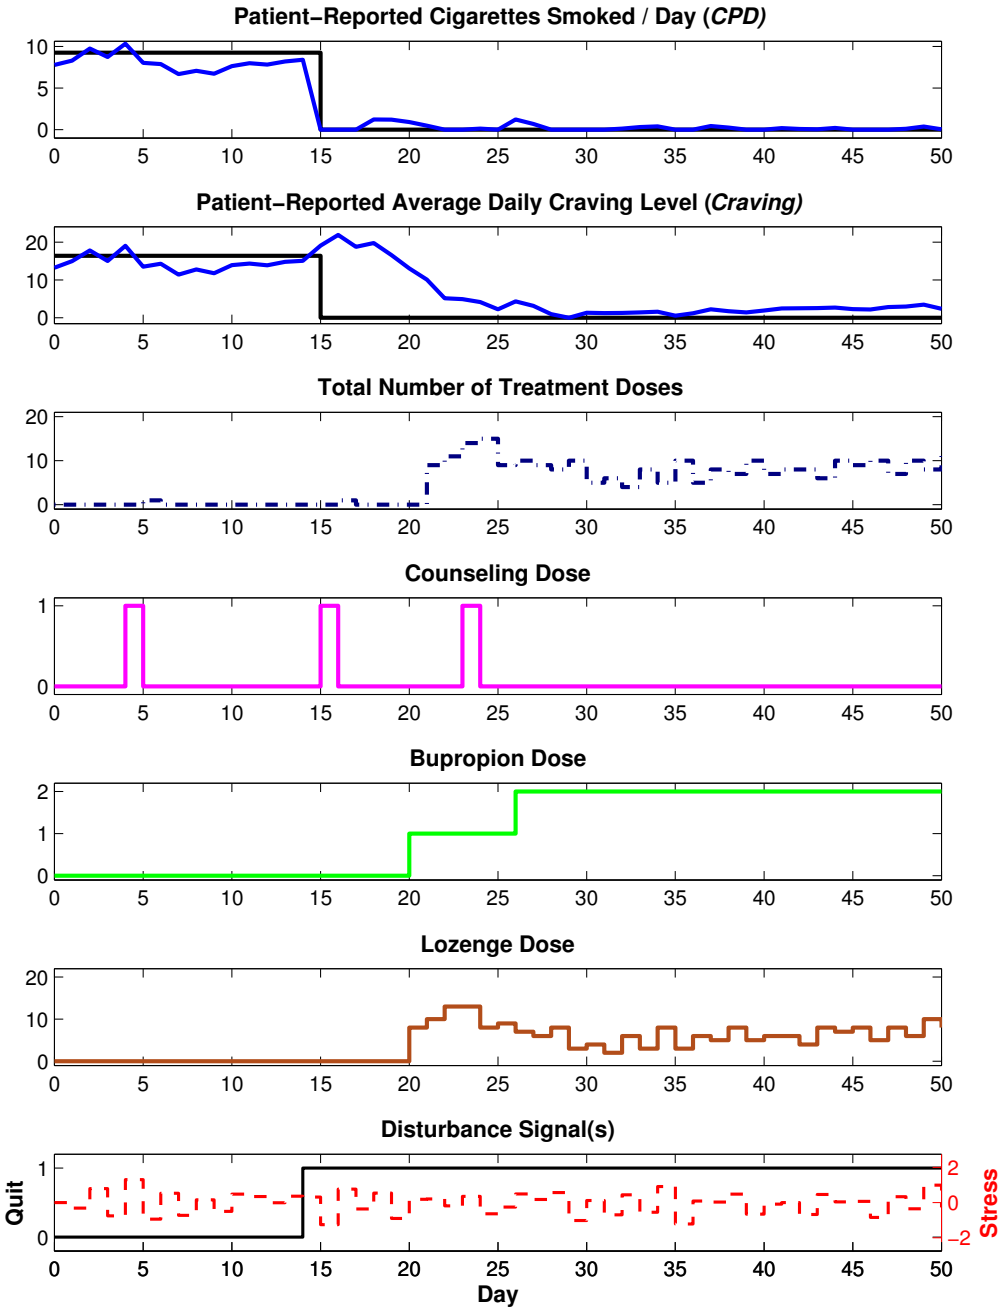

(b) Scenario 9: Outcome and Manipulated Variable Responses when a Stochastic *Stress* Disturbance is Present and $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$, $\alpha_d^S = 0.2$, $\alpha_r^{cpd} = 0$, and $\alpha_r^{crav} = 0.8$.

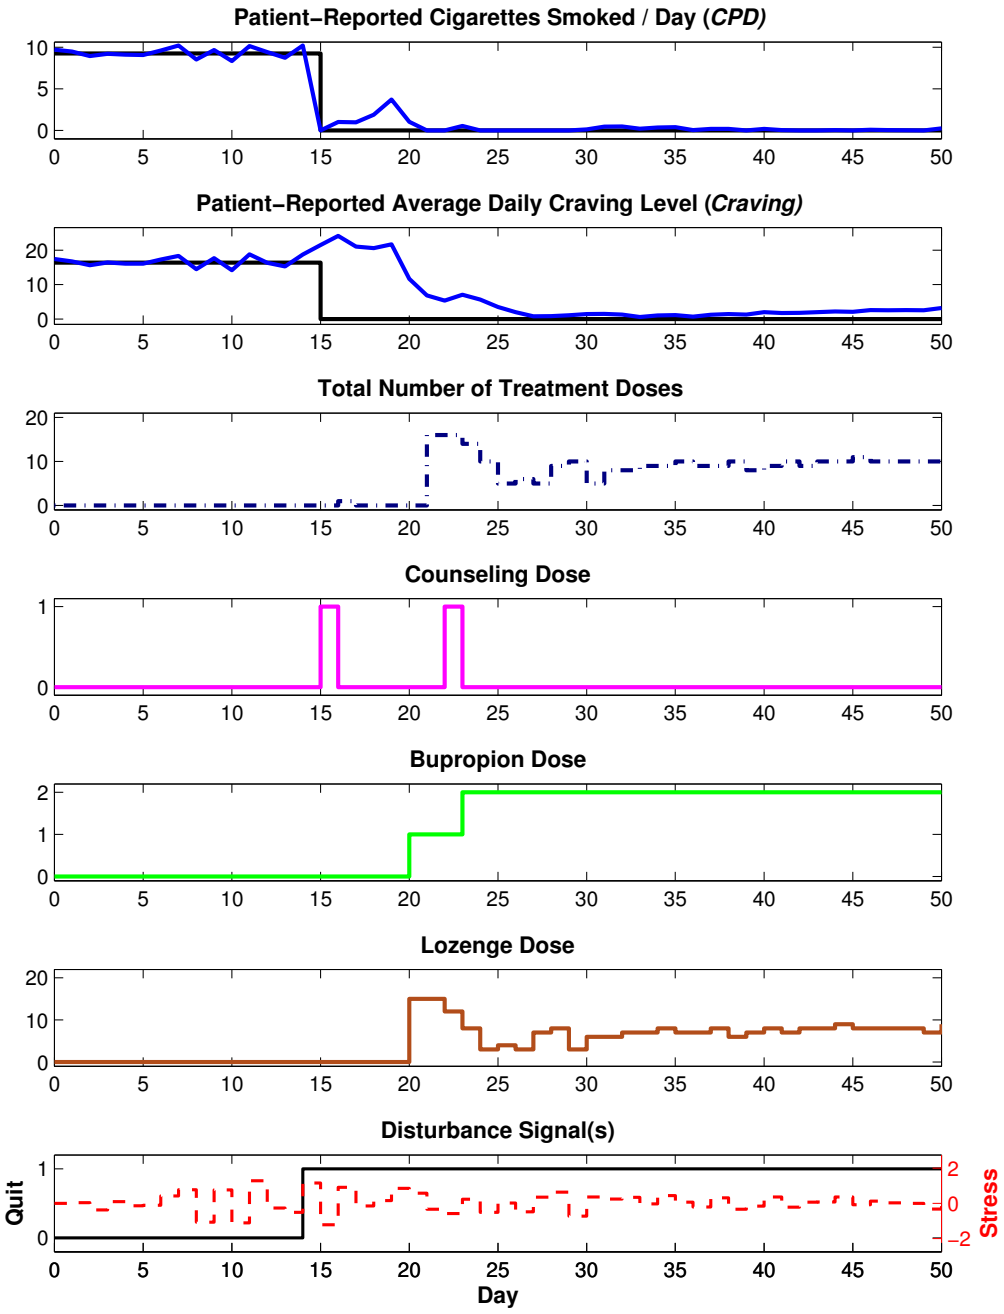

(c) Scenario 9: Outcome and Manipulated Variable Responses when a Stochastic *Stress* Disturbance is Present and $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$, $\alpha_d^S = 0.2$, and $\alpha_r^{cpd} = \alpha_r^{cpd} = 0.65$.

Figure 4.12: Scenario 9: Outcome and Manipulated Variable Responses when a Stochastic *Stress* Disturbance is Present and $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$, and $\alpha_d^S = 0.2$ for Various Combinations of α_r^{cpd} and α_r^{crav} Values.

terms of CPD and $Craving$ set point tracking in the presence of stochastic stress on average, which is achieved through reasonable dosing demands in terms of total dosing requirements, maximum daily dosing levels, and intensity of daily dosage adjustments.

4.5 Robust Performance

The previous section focused on nominal performance, i.e., when all relevant disturbances are measured and $P(s) \approx \tilde{P}(s)$. These analyses helped gain insight into the mechanics of the decision-making generally, and the sensitivity of the outcomes and dosing to certain constraints, penalty weights, and 3DoF tuning parameters. However, as previously discussed, one of the defining features of health behaviors and treatment effectiveness is patient-to-patient variability, which can enter at a multitude of biological, psychological, and environmental levels (Collins, 2006; Hamburg and Collins, 2010; Hekler *et al.*, 2013a; Kendall, 2006; Piasecki, 2006). Consequently, a critical component of this work is concerned with how effective the intervention may be in the presence of *imperfect* patient information. In other words, robust performance in this context considers how sensitive the performance of the intervention algorithm is to unmodeled perturbations (Morari and Zafiriou, 1989).

This section continues to consider the same general intervention formulation as that discussed up to this point, where the nominal models describe the representative patient (equations 3.5 through 3.8 and Table 3.2). However, three types of unmodeled perturbations are considered. Section 4.5.1 incorporates an unmeasured, stochastic influence into the measured outcomes, representing an unmeasured disturbance or measurement noise. This sort of variance is particularly apparent in the $Craving$ signals observed in the McCarthy *et al.* (2008b) clinical trial (McCarthy *et al.*, 2006). Section 4.5.2 incorporates plant-model mismatch in the form of parametric uncertainty in a subset of parameters from the representative patient's self-

regulation (*Quit*-path) and dose-response models. Section 4.5.3 incorporates plant-model mismatch by applying the intervention formulation to patients other than the representative one around whom the intervention was primarily designed. Specifically, the simulations employ $Craving(z^{-1})/Quit(z^{-1})$ and $CPD(z^{-1})/Quit(z^{-1})$ models estimated for other single subjects from the PNc group in the McCarthy *et al.* (2008b) study. These subjects' observed $Craving$ and CPD dynamics diverge from those of the representative patient. As the orders of these alternative patient models vary, Section 4.5.3 can be seen as a focused study of certain types of non-parametric uncertainty (Morari and Zafiriou, 1989).

In the presence of nontrivial, unmeasured, and unmodeled perturbations, altering the controller tuning parameters f_a^{cpd} and f_a^{crav} can significantly influence the nature of manipulated variable adjustments and controlled variable responses (see Fig. 3.5). Specifically, the speeds at which unmeasured disturbances are rejected are directly proportional to these values, where $f_a^* = 1$ (where $*$ refers to CPD or $Craving$) corresponds to the most aggressive rejection of unmeasured disturbances (Nandola and Rivera, 2013). In the following sections, these parameters will be tuned independently and in combination to examine how a clinician could influence the character of dosage assignments and intervention target tracking in the presence of unmeasured patient-to-patient variability. For simplicity, the following remain constant in the sections below unless otherwise specified: $Q_{cpd} = Q_{crav} = 10$, $Q_{wip_T} = 1$, $\alpha_r^{cpd} = \alpha_r^{crav} = \alpha_d^S = 0$, and $Stress = 0$.

4.5.1 Unmeasured, Stochastic Perturbations

Even with future advances in measurement methods (e.g., wearable technologies; Lopez-Meyer *et al.*, 2013), it will not be possible to perfectly measure or predict the effect of all time-varying factors that influence a smoker's behavior or the effect of

accurate self-reporting (such as recall bias; Raphael, 1987). Consequently, examining the robust performance of an HMPC-based cessation intervention in the context of unmodeled variance in outcome measurements is critical. In this section, the intervention performance is evaluated when the self-reported controlled variable signals are subject to measurement noise.

Significant variance is characteristic of day-to-day measurements of *Craving*, the level of which generally increases upon initiation of the quit attempt (McCarthy *et al.*, 2006). To reflect this, the following simulations subject the measured *Craving* signal to a normally distributed pseudorandom sequence, the variance of which increases as of TQD. Specifically, for $t <$ TQD, $N \sim (0, 8.95)$, and for $t \geq$ TQD, $N \sim (0, 11.90)$. The initial level of variance in *Craving* is based off of the pre-TQD variance estimated for the subject in the McCarthy *et al.* (2008b) clinical trial after whom the nominal, representative patient was patterned. The post-TQD *Craving* variance is based off of the the ratio between post-TQD and pre-TQD *Craving* variance for individual subjects from the PNc group in the McCarthy *et al.* (2008b) study who had sufficient enough data to estimate a discrete-time model; in other words, this representative subject's post-TQD *Craving* variance is nearly 30% greater than the pre-TQD period, which is reflective of the increased level of the variance generally observed in the PNc treatment condition. Fig. 4.13 depicts one realization of the stochastic signal added to the simulated, non-perturbed *Craving* responses over time. A pseudorandom sequence of the character $N \sim (0, 4.55)$ is added to simulated, non-perturbed CPD signals for the pre-TQD time period. This is based on the pre-TQD CPD variance for the clinical trial subject who acted as the basis for the representative patient; for simplicity and more straightforward performance evaluation, it is assumed that CPD is accurately reported for $t \geq$ TQD. (Anecdotal evidence suggests significant variance in CPD post-TQD is rare.)

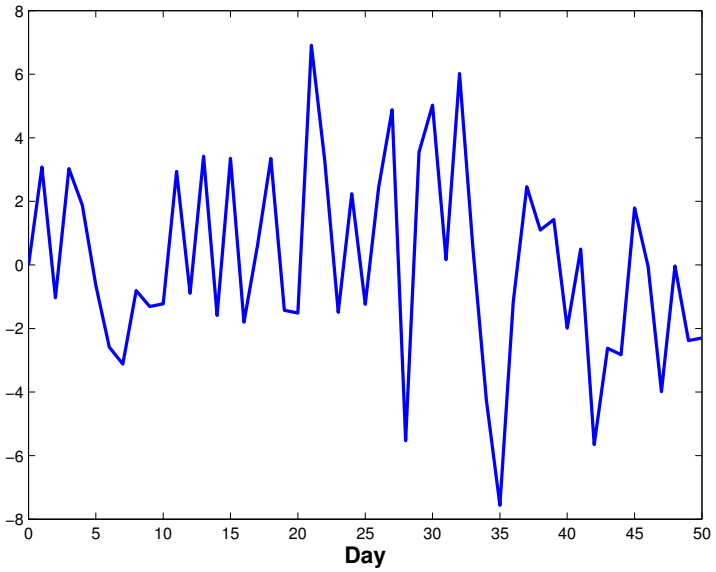

Figure 4.13: A Time Series realization of the Stochastic Signal that is Incorporated as a Measured Disturbance into the Simulated, Un-perturbed *Craving* Measurements (where TQD is Day 15).

Fig. 4.14a depicts the dosing and corresponding outcome responses for one realization of unmodeled measurement noise on *Craving* and pre-TQD *CPD* levels ($f_a^{cpd} = f_a^{crav} = 1$). Fig. 4.14a depicts considerable reduction in both *CPD* and *Craving* levels, also reflected in the relatively low offset values on day 50, 0.12 cigarettes and 4.58 points, respectively. These reductions suggest that the controller can assign effective dosing despite the presence of unmeasured, stochastic perturbations. Interestingly, the general shape of the controlled variable responses in Fig. 4.14a look similar to noise-free cases under similar conditions, per Fig. 4.5b (which did not have a $0 \leq \Delta u_b \leq 1$ constraint). Similar u_c and u_b doses are assigned in both cases, with the noisy case employing more aggressive dosage adjustments in an attempt to reject the variation in *Craving* ($var(WIP_T)$ is 12.71 versus 0.25 in the noise-free case). That said, the influence of the noise is clear in both outcomes. *Craving* approaches 50% of the baseline level four times after day 25. Due to the interrelated

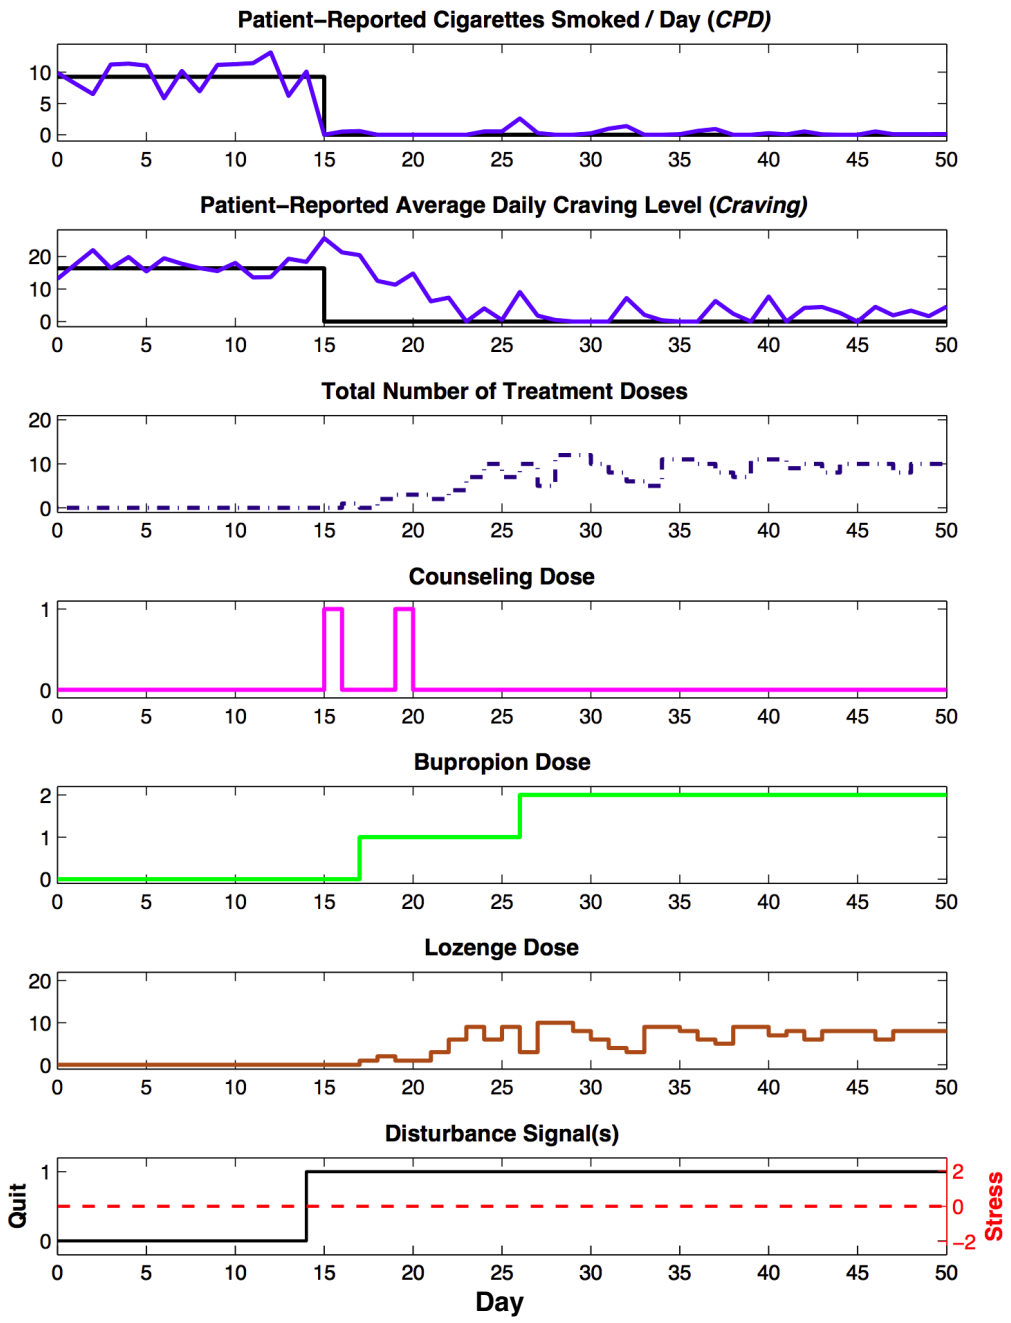

(a) Dosing and Responses when $Craving$ and CPD are Subject to Unmodeled, Stochastic Measurement Noise; $f_a^{cpd} = f_a^{crav} = 1$.

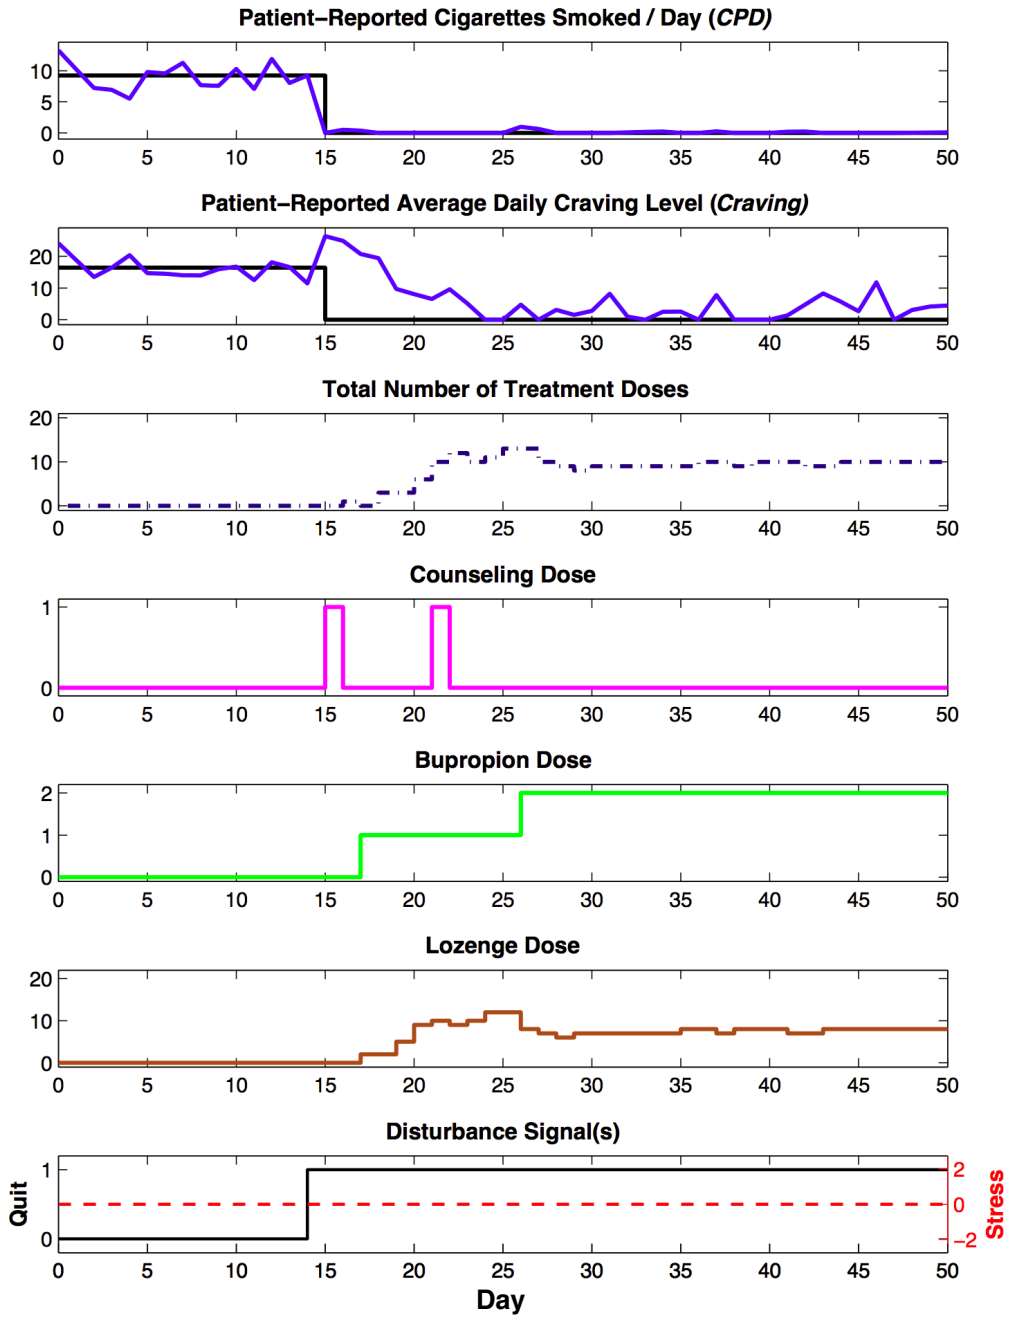

(b) Dosing and Responses when $Craving$ and CPD are Subject to Unmodeled, Stochastic Measurement Noise; $f_a^{cpd} = 0.8$ and $f_a^{crav} = 0.4$.

Figure 4.14: Dosing and Responses when $Craving$ and CPD are Subject to Unmodeled, Stochastic Measurement Noise for Various Combinations of f_a^{cpd} and f_a^{crav} Values. (See Subsection 4.5.1.)

nature of CPD and $Craving$, the dosing decisions based off of the noisy $Craving$ signal affect the CPD outcome as well, with approximately double the number of total cigarettes smoked after TQD versus the noise-free equivalent, and the maximum lapse even is 2.60 cigarettes (which was less than 1 cigarette in the case corresponding to Fig. 4.5b).

Table 4.9: Intervention Performance Metrics When CPD and $Craving$ are Subject to Measurement Noise for Various Combinations of Detuning via the Observer Gain Matrix Parameters, Averaged Across 20 Noise Realizations. (See Subsection 4.5.1.)

Parameters	f_a^{cpd}	1	1	1	0.6	0.6	0.6	0.4	0.4	0.4	0.2	0.2	0.2
	f_a^{crav}	1	0.6	0.2	1	0.6	0.2	1	0.6	0.2	1	0.6	0.2
Performance, $t \geq$ TQD	t_r^{cpd}	0	0	0	0	0	0	0	0	0	0	0	0
	t_r^{crav}	8.9	7.95	9.05	8.75	8.85	9.25	8.85	9.35	8.75	9.65	9	9.75
	e_{50}^{cpd}	0.17	0.1	0.06	0.35	0.16	0.04	0.17	0.08	0.05	0.24	0.06	0.09
	e_{50}^{crav}	3.7	4.18	2.36	4.23	4.77	4.56	3.43	3.77	3.35	3.42	3.54	3.78
	Total Cigs	9.45	6.39	4.77	6.97	4.61	1.35	4.85	2.22	2.09	4.89	2.38	1.35
	Days CPD=0	14.6	18.2	19	17.45	19.25	27.8	22	25.7	29.3	22.6	26.65	31.45
	$max(cpd)$	1.75	1.29	1.17	1.3	0.75	0.32	1.04	0.58	0.38	1	0.65	0.48
	$max(crav)$	24.01	21.18	22.71	24.2	24.81	21.59	23.08	24.43	22.57	23.86	24.9	24.38
	$max(WIP_T)$	15.25	16.45	15.25	13.9	14.55	13.8	14.2	13.5	12.55	14.35	14.95	15.15
	J_e^{cpd}	9.36	6	4.46	5.22	2.51	0.66	3.84	1.03	1.44	3.07	1.25	0.82
	J_e^{crav}	2622	2054	2323	2367	2370	2022	2138	2162	1997	2321	2222	2128
	$J_e^{wip_T}$	2624	3090	2739	2475	3120	2914	2761	2716	2779	2817	3079	3591
	$var(cpd)$	0.19	0.13	0.1	0.11	0.05	0.01	0.08	0.02	0.02	0.07	0.03	0.02
	$var(crav)$	42.95	34.78	40.59	40	40.57	34.6	36.56	38.27	35.57	39.85	39.72	39.63
	$var(WIP_T)$	14.3	15.51	14.35	10.71	12.08	10.59	11.23	10.47	10.76	11.53	12.06	13.65
	$J_I^{u_c}$	2.15	2.45	2.25	2.25	2.15	2.25	2.3	2.25	2.65	2.2	2.3	2.5
	$J_I^{u_b}$	112	110	111	113	109	118	112	117	109	114	113	105
	$J_I^{u_l}$	1766	2131	1883	1629	2177	1918	1896	1804	1906	1926	2144	2664
	$var(u_c)$	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.07	0.06	0.06	0.07
	$var(u_b)$	0.4	0.35	0.44	0.34	0.35	0.3	0.31	0.33	0.32	0.33	0.33	0.31
	$var(u_l)$	11.33	12.22	10.81	7.84	8.58	7.14	8.63	7.33	7.71	8.81	9.04	10.3

The plots in Fig. 4.14a correspond to the most aggressive possible dosing for the penalty weight combination, i.e., no detuning is present as specified by $f_a^{cpd} = f_a^{crav} = 1$. Fig. 4.15 contains a surface plot of the J_e^y values relative to the $f_a^{cpd} =$

$f_a^{crav} = 1$ case, averaged across 20 realizations of the noise, as a function of the tuning parameters in the observer. A selection of the corresponding performance metrics are found in Table 4.9. The average, scaled variance of the WIP_T construct is also depicted. Fig. 4.15c reflects that detuning via the observer gain has a more complex effect on the total dosing levels than has been observed up to this point. That said, Fig. 4.15a indicates that filtering conditions that suppress tracking of the *Craving* noise leads to more favorable tracking of the CPD target post-TQD. Fig. 4.15d suggests that adjusting the f_a^{cpd} parameter has a greater effect on the intensity with which the treatments are adjusted, which appears to correspond to specific dosing combinations and sequences that facilitate more consistent tracking of the CPD target. The decreased amount of dosing adjustments and the improved tracking of the CPD target suggests clinical implementation of an intervention of this sort would benefit from a tuning scenario in which $f_a^{cpd} < 1$ and $f_a^{crav} < 1$. Fig. 4.14b depicts an example, where $f_a^{cpd} = 0.8$ and $f_a^{crav} = 0.4$. Comparing this figure to Fig. 4.14a indicates that the detuning makes dosage adjustments less sensitive to the *Craving* noise, particularly adjustments in u_l, as expected. This more patient-friendly dosing also leads to improved CPD performance versus the $f_a^{cpd} = f_a^{crav} = 1$ case, reflected in the lower amounts of total smoking and lower peak lapse events during the quit attempt. In other words, CPD performance is also less sensitive to noise in *Craving* when the controller is detuned in order to discourage aggressive noise rejection.

4.5.2 Plant-Model Mismatch due to Parametric Uncertainty

This section considers robust performance in the presence of plant-model mismatch due to parametric uncertainty. (Note, the remaining analyses do not incorporate unmodeled noise in the CPD and *Craving* measurements.) Specifically, param-

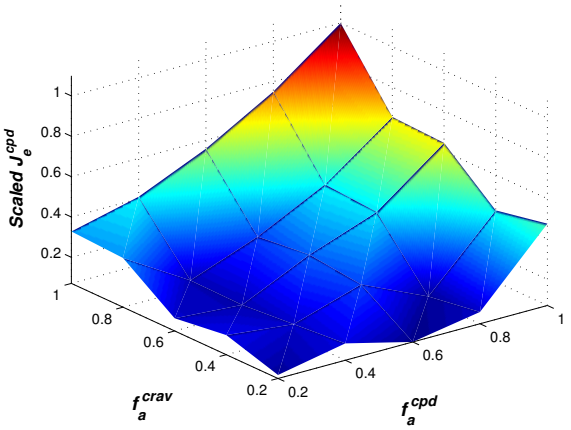

(a) Relative J_e^{cpd} Values.

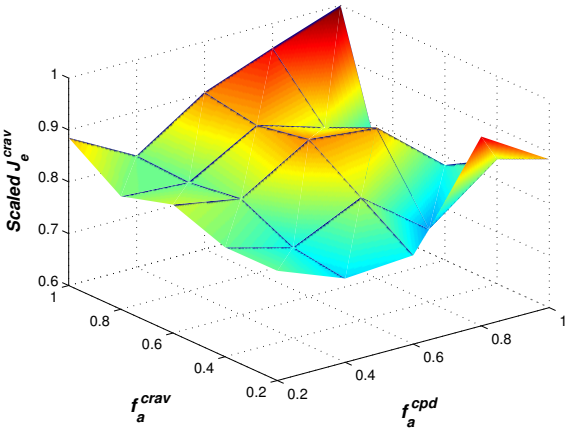

(b) Relative J_e^{crav} Values.

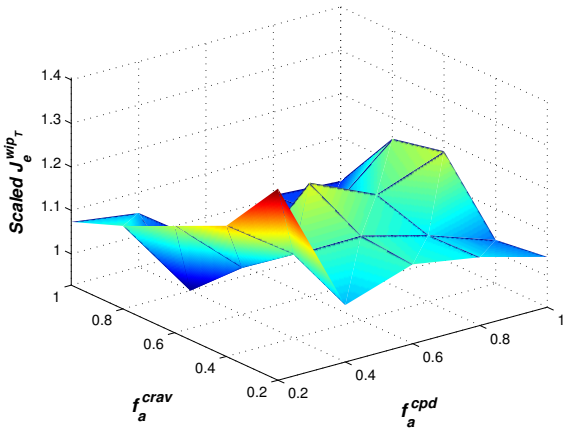

(c) Relative $J_e^{wip_T}$ Values.

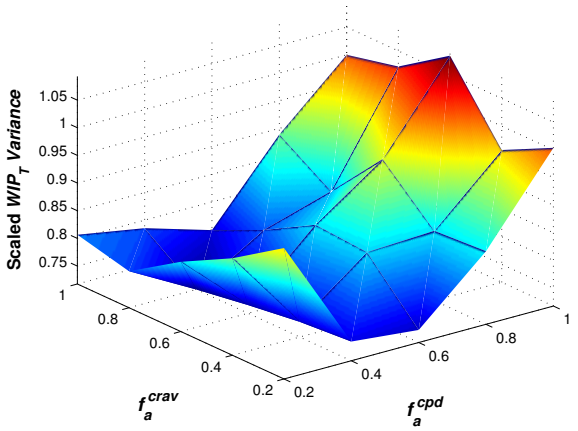

(d) Relative $var(WIP_T)$ Values.

Figure 4.15: Surface Plots of Intervention Performance Metrics for Different Levels of Detuning via f_a^y, as Averaged Across 20 Realizations of Noise and Scaled Relative to the $f_a^{cpd} = f_a^{crav} = 1$ Values. (See Subsection 4.5.1.)

eters from the self-regulation models—the open-loop *Quit*-path models—are altered to indicate that the plant resumes smoking more quickly and/or experiences greater peak *Craving* levels than is captured in the nominal models; parameters in the open-loop dose-response models are also altered to represent uncertainties that lead to less effective treatment components. As exogenous, unmeasured disturbances and plant-model mismatch factor in to control decisions in an equivalent manner, this section also briefly examines how adjusting the tuning parameters in the observer gain matrix affect the sensitivity of control action to this mismatch. As the dose-response models draw from literature and theory more than data, and as these analyses are limited to simulation, the focus of this section is on evaluating the effects of these uncertainties generally, and therefore each source of uncertainty is introduced and examined independently. Tables 4.10 and 4.11 contain the performance metrics for various levels of parametric uncertainty for the mismatch cases examined; Figures 4.16 and 4.17 depict the dosing demands and responses for a subset of these cases.

Uncertainty in the *Quit*-Path Models

Analyzed in Chapter 2, the responses of CPD and $Craving$ to initiation of a quit attempt—when represented by the closed-loop transfer functions in equations 2.35 and 2.36 and the corresponding structures in equations 2.30 through 2.32—are sensitive to moderate to large changes in the K_d and K_1 values. Specifically, decreases in the K_d magnitude leads to smaller decreases in CPD on TQD, indicating an inability of the patient to fully quit on their own at any point. This also leads to smaller peak $Craving$ magnitudes. Increasing the magnitude of K_1 serves to increase the speed at which smoking is fully resumed, increase the CPD overshoot seen after TQD during relapse, and increase peak $Craving$ levels. Table 4.10 contains the performance metrics for the following cases in which K_d and K_1 are altered independently: the magnitude of K_d is 30% and 60% smaller than baseline smoking levels and the magnitude of K_1 increases by 50% and 100%. Fig. 4.16a and Fig. 4.16c depict the dosing demands and responses for the latter of each of these mismatch scenarios ($f_a^{cpd} = f_a^{crav} = 1$); the open-loop CPD and $Craving$ responses of the uncertainty models are also depicted (dashed grey).

Compared to the equivalent nominal performance case, Fig. 4.16a reflects that uncertainty in K_d primarily influences dosing in the first two weeks of a quit attempt; early aggressive u_l dosing serves to bring CPD down to target levels, with dosing between days 20 and 30 more hesitant than the corresponding nominal performance case (e.g., Fig. 4.5b). This is due to the lower peak in $Craving$ on TQD that results from decreased K_d magnitudes, meaning less aggressive dosing is required over this time interval to bring $Craving$ toward its target. Fig. 4.16a and Table 4.10 suggest that intervention performance is not significantly degraded by decreased K_d magnitudes; while smoking abstinence is not achieved until after TQD (reflected in the

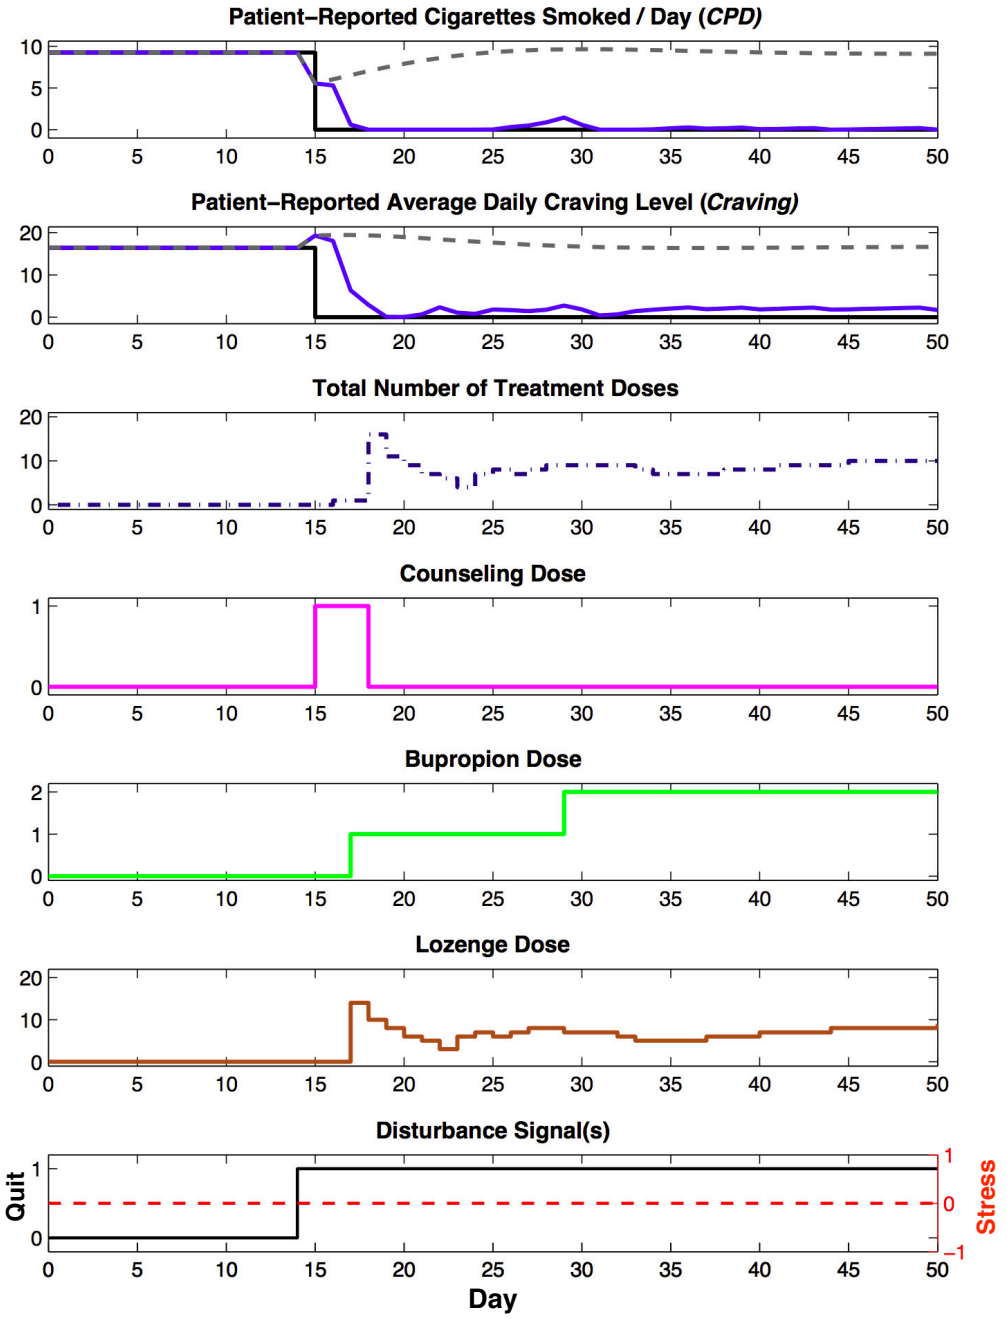

(a) Dosing and Responses (Solid Blue) in the Presence of Plant-Model Mismatch where the K_d in the Plant Models is 40% of the Representative Subject's Baseline CPD Level; $f_a^{cpd} = f_a^{crav} = 1$. Open-loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

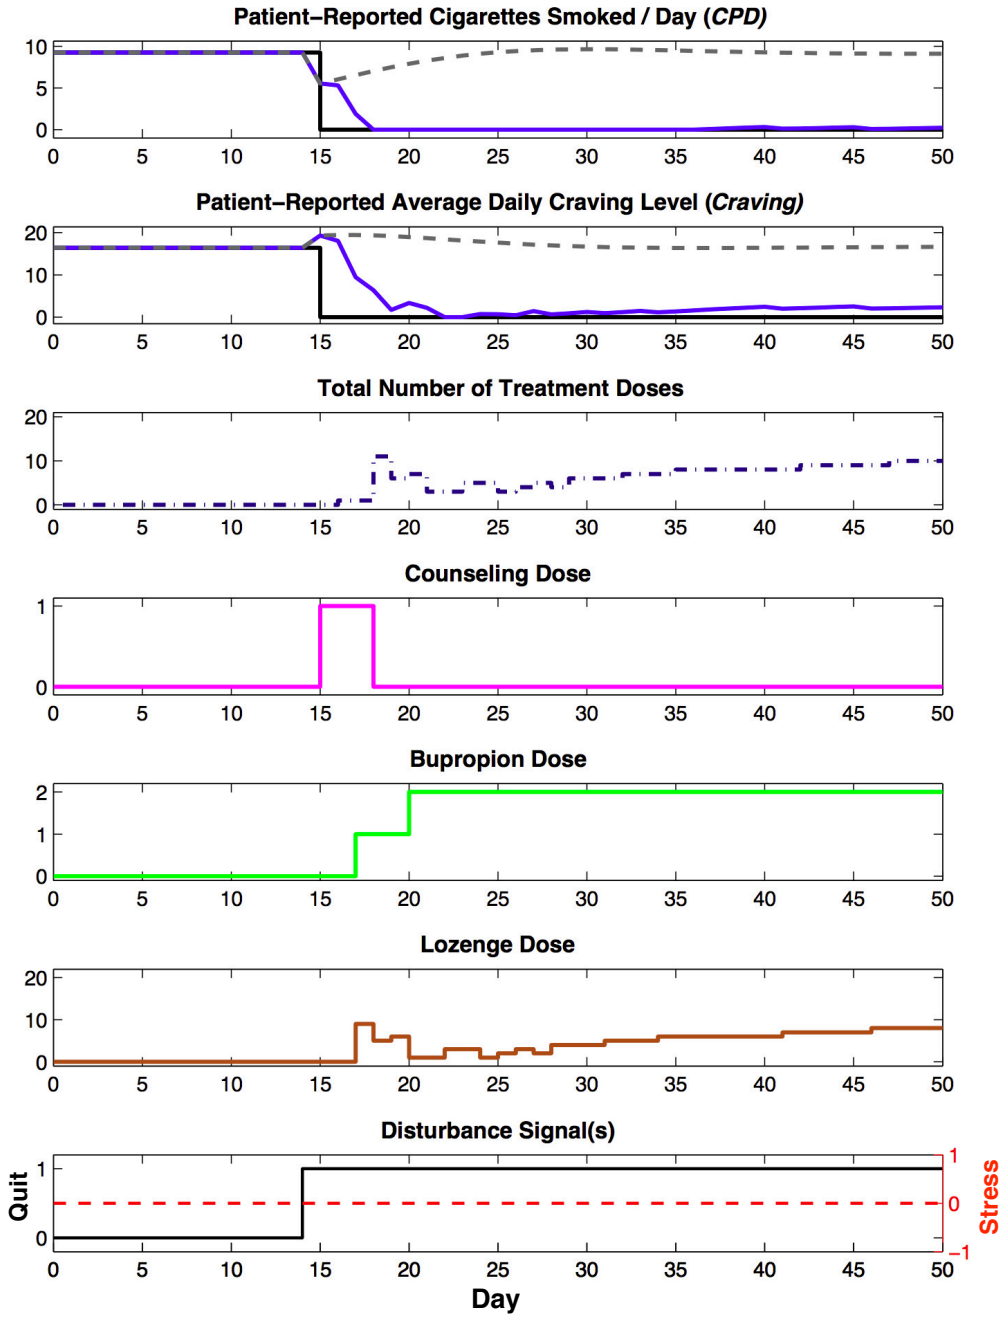

(b) Dosing and Responses (Solid Blue) in the Presence of Plant-Model Mismatch where the K_d in the Plant Models is 40% of the Representative Subject's Baseline CPD Level; $f_a^{cpd} = f_a^{crav} = 0.3$. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

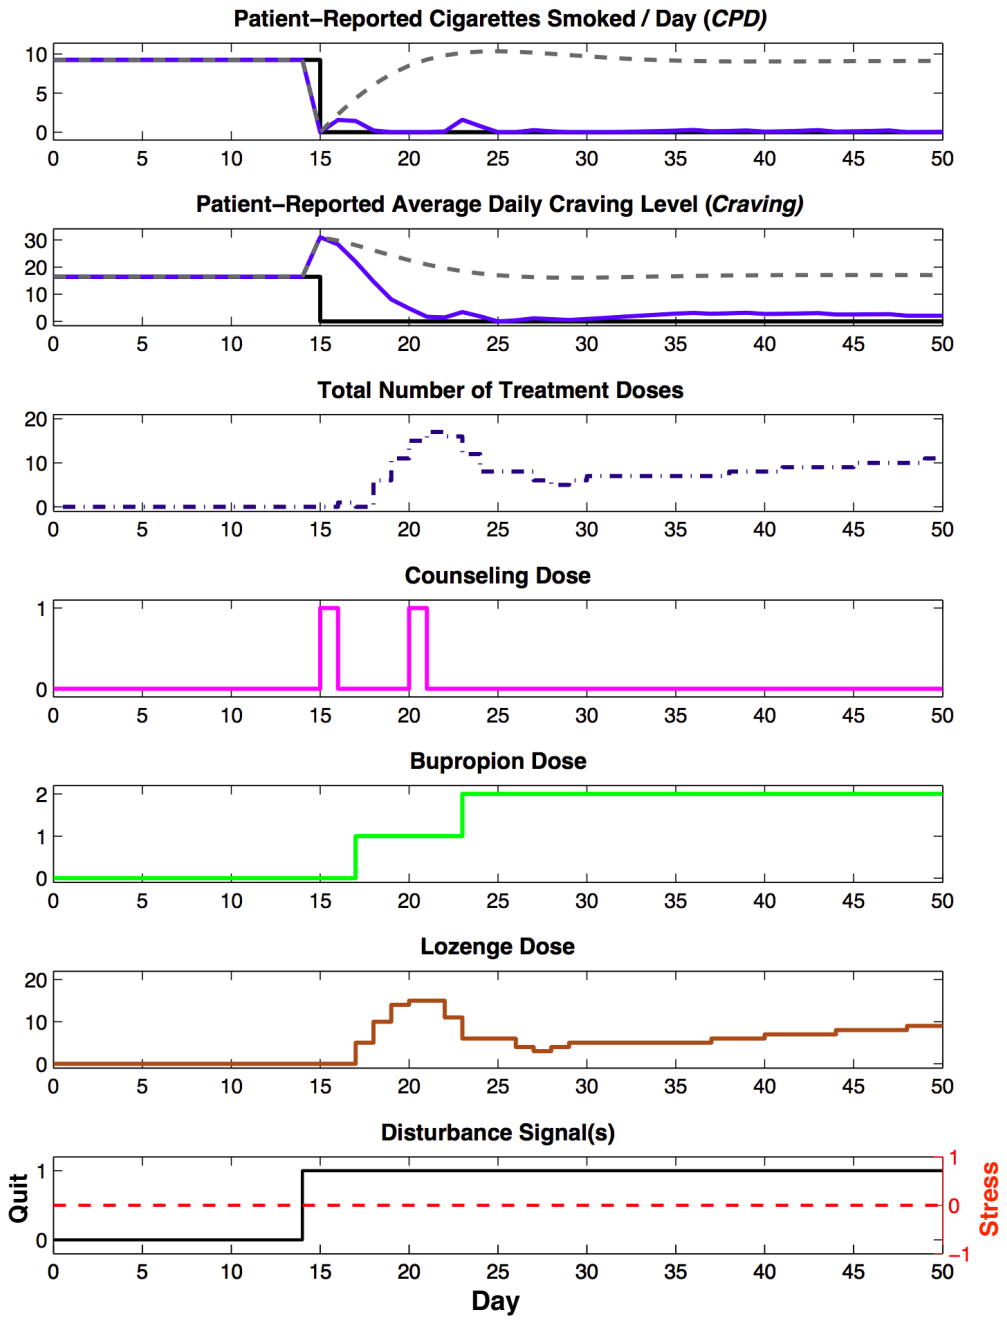

(c) Dosing and Responses (Solid Blue) in the Presence of Plant-Model Mismatch where the K_1 in the Plant Models is 100% Greater than that for the Representative Subject; $f_a^{cpd} = f_a^{crav} = 1$. Open-loop Responses of CPD and $Craving$ to $Quit$ are also Depicted (Dashed Grey).

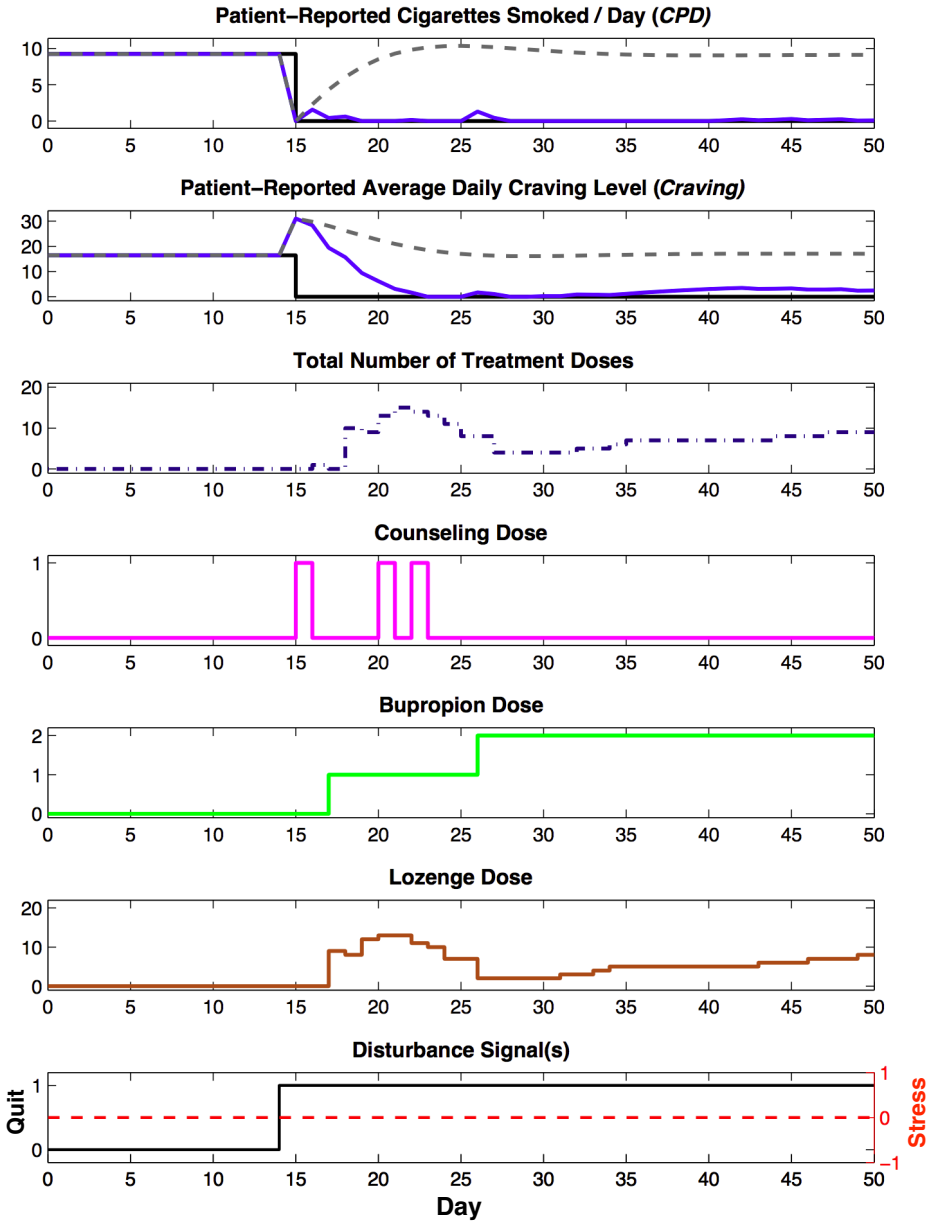

(d) Dosing and Responses (Solid Blue) in the Presence of Plant-Model Mismatch where the K_1 in the Plant Models is 100% Greater than that for the Representative Subject; $f_a^{cpd} = f_a^{crav} = 0.3$. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

Figure 4.16: Dosing and Responses (Solid Blue) in the Presence of Plant-Model Mismatch Introduced into Individual Parameters in Equations 2.35 and 2.36. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey). (See Subsection 4.5.2.)

Table 4.10: Intervention Performance Metrics in the Presence of Uncertainty in Two Parameters of the *Quit*-Path Nominal Models, Where the Specified Uncertainty in the Parameter is Relative to the Value of the Parameter in the Representative Subject Models (See Subsection 4.5.2).

Uncertain parameter		K_d	K_d	K_d	K_d	K_1	K_1	K_1	K_1
Degree of uncertainty		70%	70%	40%	40%	150%	150%	200%	200%
Tuning parameter	f_a^{cpd}	1	0.3	1	0.3	1	0.3	1	0.3
	f_a^{crav}	1	0.3	1	0.3	1	0.3	1	0.3
Performance, $t \geq$ TQD	t_r^{cpd}	3	2	3	3	0	0	0	0
	t_r^{crav}	11	9	4	7	12	10	10	8
	e_{50}^{cpd}	0.08	0.04	0.00	0.24	0.23	0.01	0.06	0.10
	e_{50}^{crav}	2.42	2.35	1.68	2.31	3.07	2.50	2.12	2.47
	Total Cigs	8.45	7.15	17.32	15.39	4.35	3.87	8.67	6.23
	Days CPD=0	16	22	12	19	16	24	9	20
	$max(cpd)$	2.90	2.90	5.55	5.55	1.09	1.05	1.59	1.57
	$max(crav)$	21.51	21.51	19.32	19.32	27.34	27.34	30.99	30.99
	$max(WIP_T)$	10	10	16	11	14	17	17	15
	J_e^{cpd}	17.13	16.41	63.27	63.07	2.71	2.63	8.32	5.31
	J_e^{crav}	1409.42	1363.07	848.63	931.81	2536.92	2271.14	2709.35	2659.33
	$J_e^{wip_T}$	1840	1805	2568	1850	1794	2546	2911	2332
	$var(cpd)$	0.43	0.43	1.57	1.61	0.06	0.06	0.18	0.12
	$var(crav)$	25.62	26.18	16.44	17.92	47.18	45.94	53.90	54.60
	$var(WIP_T)$	6.86	5.48	8.45	7.90	12.14	14.03	13.16	12.15
	$J_I^{u_c}$	3	3	3	3	3	3	2	3
	$J_I^{u_b}$	118	127	100	127	118	109	118	109
	$J_I^{u_l}$	1121	1069	1780	1120	1162	1782	2044	1608
	$var(u_c)$	0.08	0.08	0.08	0.08	0.08	0.08	0.05	0.08
	$var(u_b)$	0.32	0.28	0.37	0.28	0.32	0.35	0.32	0.35
	$var(u_l)$	4.87	3.96	5.90	6.29	11.33	12.50	11.15	10.94

nonzero t_r^{cpd} values), lapses after day 20 appear manageable even when the K_d magnitude is only 40% of the baseline CPD level, and there may be improved $Craving$ performance overall (e.g., lower t_r^{crav} values than in previous scenarios). Good—even improved—CPD performance can also be obtained with $f_a^{cpd} < 1$ and/or $f_a^{crav} < 1$ (e.g., Fig. 4.16b). In this case, slowing the speed at which the mismatch is rejected leads to a suppression of aggressive changes in u_l early in the quit attempt. Although this leads to a slightly increased post-TQD period on nontrivial smoking, it also leads

to earlier implementation of the second u_b dose, which helps bring CPD to target levels relatively quickly and bring down $Craving$ quickly as well.

Examining the cases considered here with uncertainty in K_1 (see equations 2.30, 2.35, and 2.36) suggests that the intervention can be efficacious overall despite the mismatch, but at the expense of some post-TQD lapse events and aggressive dosing early within the quit attempt (when $f_a^{cpd} = f_a^{crav} = 1$). This is seen with Fig. 4.16c. However, even relatively aggressive detuning via the observer gain has relatively minor effects on that performance and the corresponding dosing demands, as apparent in Fig. 4.16d.

Altogether, longer-term intervention efficacy appears to be relatively robust to uncertainty in the degree to which a patient can quit on their own on TQD (i.e., smaller values of K_d relative to baseline CPD levels) and the K_1 parameter (which primarily affects the open-loop increase in $Craving$), however uncertainty in the latter scenario is managed at the expense of aggressive dosing of all treatment components early in the intervention, which is not well-handled by decreases in f_a^* values. Practically, this suggests that intervention performance may be most sensitive to uncertainty in K_1, as the aggressive dosing assigned by the intervention in this case is less patient-friendly than would be preferred; this likely increases the probability of non-adherence, ultimately leading to degraded performance.

Uncertainty in the Dose-Response Models

Table 4.11 contains the performance metrics for the cases in which there is uncertainty in one parameter in the dose-response models; specifically, 20% and 40% reductions in the gain of each open-loop, discrete-time dose-response model. Generally, with $f_a^{cpd} = f_a^{crav} = 1$, each scenario with these independent uncertainties does lead to CPD and $Craving$ values that are at or near target levels by day 50, and CPD remains within

5 cigarettes of its target over the entire time period. Uncertainty in the counseling models appears to have the least effect on dosing and outcome performance, as could be expected given the limited use of counseling overall. The effect of uncertainty in the gain of the lozenge model appears to be most consequential to the lozenge assignments, generally leading to increased amounts of lozenge assignment at most time points compared to most previous cases. Intervention dosing and efficacy appears to be most sensitive to uncertainty in the effect of bupropion. Decreasing the f_a^* values under these uncertainty conditions appears to have some minor effect on u_l dosing but relatively small changes in performance. As seen in Fig. 4.17a, for $f_a^* = 1$, the lower gain in the bupropion dose-response model leads to increased counseling and generally a sustained increase in daily lozenge dose and adjustment frequency; this still leads to larger amounts lapse and fewer cigarette-free days (see Table 4.11). For the independent, 40% mismatch in the gains of the counseling and lozenge dose-response gains, detuning such that $f_a^{cpd} = f_a^{crav} = 0.3$ leads to more patient-friendly dosing in that the frequency of u_l adjustment is decreased overall. However, for the equivalent mismatch in bupropion dose-response models, total amount of lozenges assigned decreased slightly, bringing in the maximum bupropion dose earlier (see Fig. 4.17b).

While the outcome targets are obtained in these simulations, they each incorporate only one distinct source of parameter uncertainty. It may suggest that the character of dosing schedules and lapse events will be more sensitive to a confluence of uncertainties of multiple dose-response models, as opposed to the *Quit*-path models.

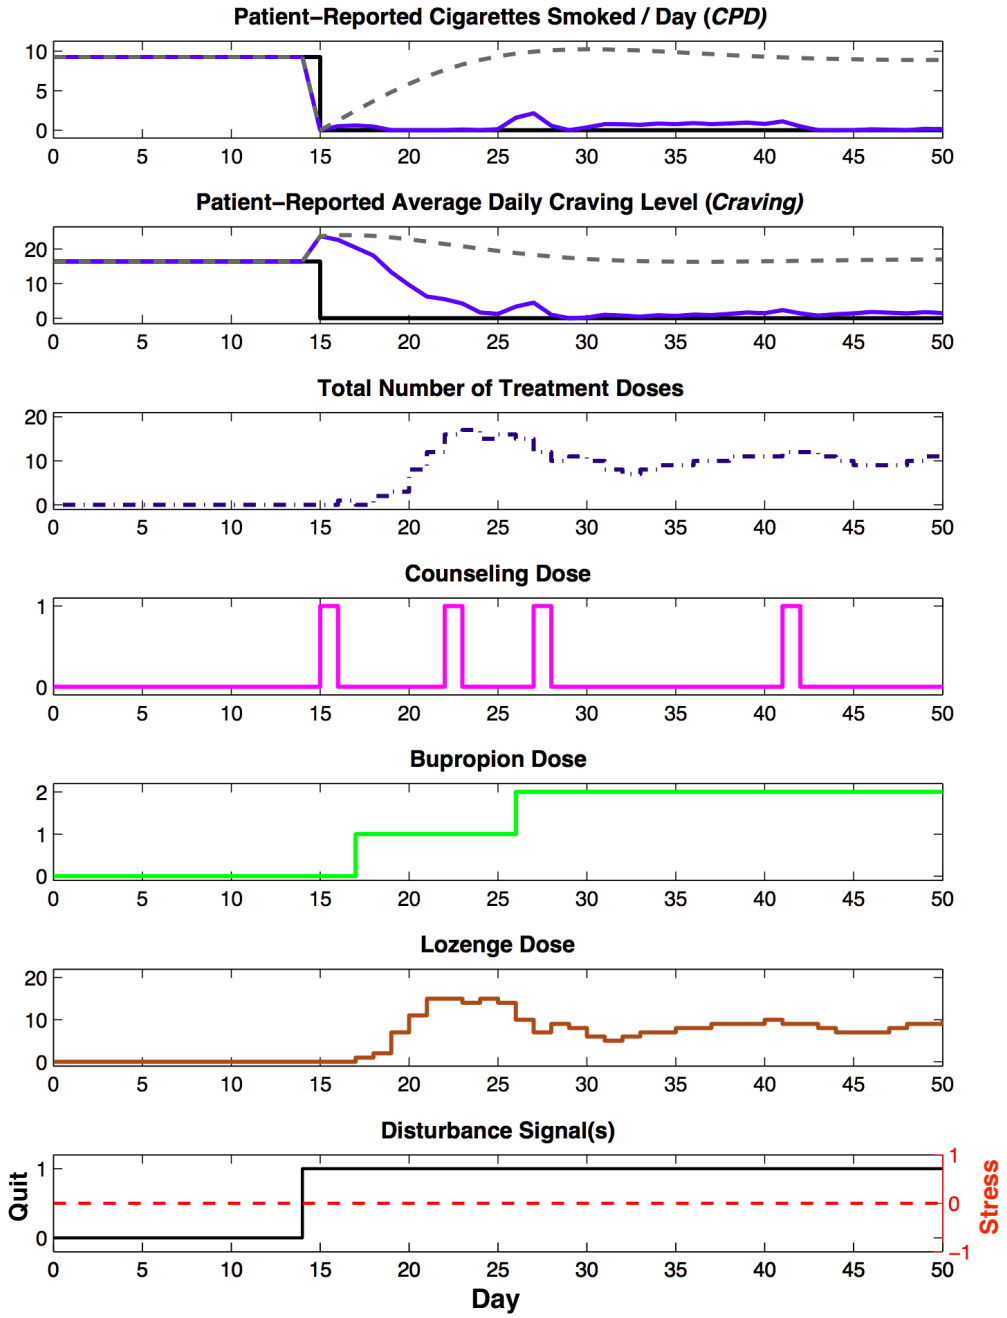

(a) Dosing and Responses where the Gain Values in the Bupropion-Response Models for the Plant are 40% of the Corresponding Gain Values for the Nominal Models; $f_a^{cpd} = f_a^{crav} = 1$. Open-loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

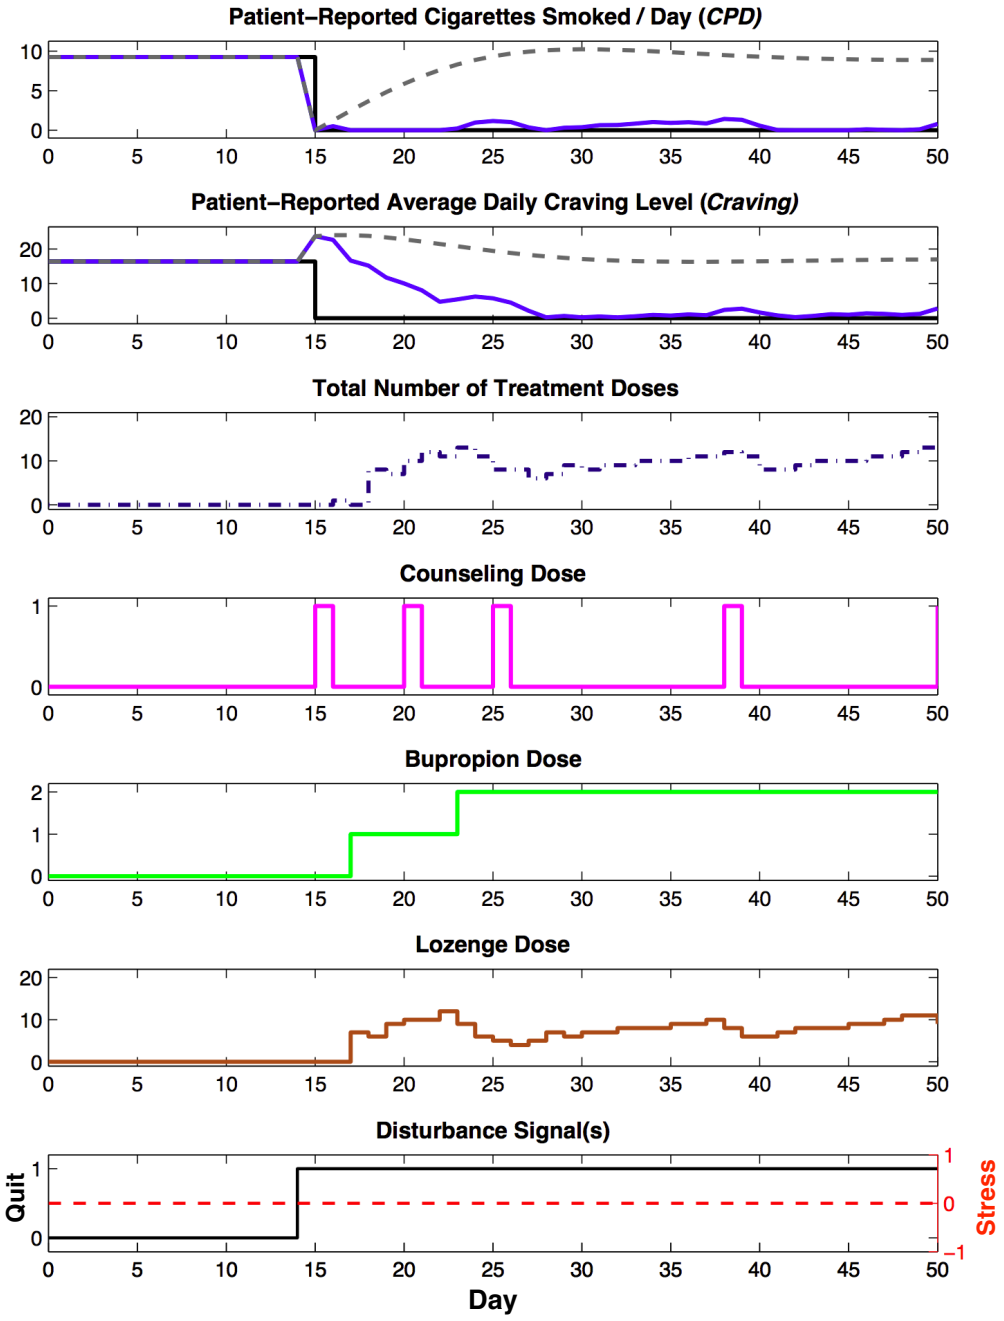

(b) Dosing and Responses where the Gain Values in the Bupropion-Response Models for the Plant are 40% of the Corresponding Gain Values for the Nominal Models; $f_a^{cpd} = f_a^{crav} = 0.3$. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

Figure 4.17: Dosing and Responses where the Gain Values in the Bupropion-Response Models for the Plant are 40% of the Corresponding Gain Values for the Nominal Models. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey). (See Subsection 4.5.2.)

Table 4.11: Intervention Performance Metrics in the Presence of Uncertainty in the Gain Parameters of the Discrete-Time Dose-Response Models, Where the Degree of Uncertainty is Relative to the Gain for the Equivalent Nominal Model (See Subsection 4.5.2).

Dose-response model		Counseling	Counseling	Counseling	Bupropion	Bupropion	Bupropion	Lozenge	Lozenge	Lozenge
Degree of uncertainty		80%	60%	60%	80%	60%	60%	80%	60%	60%
Tuning parameters	f_a^{cpd}	1	1	0.3	1	1	0.3	1	1	0.3
	f_a^{crav}	1	1	0.3	1	1	0.3	1	1	0.3
Performance, $t \geq$ TQD	t_r^{cpd}	0	0	0	0	0	0	0	0	0
	t_r^{crav}	12	15	12	15	14	13	13	13	12
	e_{50}^{cpd}	0.15	0.01	0.06	0.04	0.12	0.78	0.10	0.11	0.12
	e_{50}^{crav}	2.85	2.63	2.71	1.64	1.53	2.83	2.78	3.18	2.86
	Total Cigs	7.19	4.31	6.28	6.62	16.43	15.14	7.10	5.46	6.99
	Days CPD=0	10.00	16.00	24.00	12.00	11.00	13.00	11.00	15.00	12.00
	$max(cpd)$	0.92	0.99	1.54	1.24	2.15	1.42	0.90	1.08	0.88
	$max(crav)$	23.70	23.70	23.70	23.70	23.70	23.70	23.70	23.70	23.70
	$max(WIP_T)$	14.00	15.00	16.00	16.00	17.00	13.00	15.00	18.00	19.00
	J_e^{cpd}	3.31	2.66	6.17	4.20	16.33	13.78	3.24	3.05	2.88
	J_e^{crav}	2322.52	2339.46	2075.02	2226.66	2249.13	2074.67	2447.69	2534.90	2279.65
	$J_e^{wip_T}$	2984.00	2820.00	3884.00	3453.00	3934.00	3329.00	3733.00	3380.00	5318.00
	$var(cpd)$	0.05	0.06	0.14	0.09	0.25	0.21	0.05	0.06	0.04
	$var(crav)$	41.22	41.74	39.42	43.34	43.58	38.45	42.73	42.29	38.05
	$var(WIP_T)$	9.97	10.09	11.51	16.88	17.39	10.25	14.37	18.32	17.23
	$J_I^{u_c}$	2.00	3.00	3.00	3.00	4.00	5.00	2.00	3.00	2.00
	$J_I^{u_b}$	118.00	118.00	100.00	109.00	109.00	118.00	109.00	109.00	118.00
	$J_I^{u_l}$	2014.00	1885.00	2865.00	2548.00	2901.00	2292.00	2712.00	2501.00	4073.00
	$var(u_c)$	0.05	0.08	0.08	0.08	0.10	0.12	0.05	0.08	0.05
	$var(u_b)$	0.32	0.32	0.37	0.35	0.35	0.32	0.35	0.35	0.32
	$var(u_l)$	6.34	6.99	8.03	14.08	13.82	6.77	9.82	15.72	11.79

4.5.3 Plant-Model Mismatch due to Non-Parametric Uncertainty in the Open-Loop Cessation Models

Section 4.5.2 examined uncertainty in two individual parameters of the open-loop *Quit* disturbance models. Now, uncertainty in the *Quit*-path disturbance models is non-parametric. Specifically, black-box, SISO, discrete-time models of various lag orders were estimated for four subjects from the PNc group in the McCarthy *et al.* (2008b) study, other than that after whom the representative patient is based. Con-

Table 4.12: Structures and Parameter Estimates of the Alternate Patients Considered, Estimated as SISO ARX Models from Four Single Subject Data Sets in the PNc Group of the McCarthy *et al.* (2008b) Study (See Subsection 4.5.3).

Subject	**Outcome**	**Baseline**	$[n_a, n_b, n_k]$	$\{a_1, \ldots, a_{n_a}\}$	$\{b_1, \ldots, b_{n_b}\}$
1	*CPD*	10.6	[8 4 2]	{-0.470, -0.1978, 0.043, 0.033, -0.239, 0.082, 0.077, 0.080}	{-3.5019, -2.9119, 2.3128, 2.4737}
1	*Craving*	25.6	[8 8 1]	{-0.364, 0.080, -0.109, -0.176, -0.224, -0.034, -0.133, -0.181}	{4.460, -10.567, 13.038, -2.965, -0.123, 0.109, -2.046, -2.998}
2	*CPD*	30.0	[1 3 1]	-0.583	{2.169, 10.420, -7.749}
2	*Craving*	17.2	[3 1 1]	{-0.472, 0.185, -0.011}	-20.140
3	*CPD*	24.2	[1 3 1]	-0.201	{3.160, 2.860}
3	*Craving*	30.8	[1 2 1]	0.133	{2.172, -26.130, 26.190}
4	*CPD*	20.3	[1 2 1]	-0.158	{-20.330, 3.316}
4	*Craving*	23.8	[4 2 1]	{-0.507, -0.204, 0.091, -0.194}	{1.955, -5.174}

sidering these alternate models as the plant's open-loop cessation dynamics facilitates analysis of robust performance where plant-model mismatch resulting from a confluence of parameter uncertainties and unmodeled complexities, such as nonlinearities that are likely characteristic of patient-to-patient variability. Generally, these analyses also help assess how appropriate the representative patient is as a basis for design of an intervention of this nature.

Table 4.12 contains the orders of the ARX structures, parameter estimates, and average pre-TQD levels of *Craving* and *CPD* for the four subjects. These sets of dynamics are examined here as they represent several classes of experimentally-observed dynamic patterns. Subject 1 nearly quits on their own initially, but eventually resumes smoking to a degree, and also features more complex *Craving* dynamics. Subject 2 nearly abstains on their own, but increased *Craving* levels are sustained for several weeks into the quit attempt. Subject 3 is less successful than the representative patient in their independent quit attempt, unable to reduce smoking for more than one day (not on TQD), which corresponds to a sustained increase in *Craving*.

Table 4.13: Intervention Performance Metrics in the Presence of Uncertainty in the Gain Parameters of the Discrete-Time Dose-Response Models, Where the Degree of Uncertainty is Relative to the Gain for the Equivalent Nominal Model (See Subsection 4.5.3).

Subject		1	1	2	2	2	3	3	4	4	4
Tuning parameters	α_r^{cpd}	0	0	0	0.8	0	0	0	0	0.8	0
	α_r^{crav}	0	0	0	0.8	0	0	0	0	0.8	0
	f_a^{cpd}	1	0.2	1	1	0.2	1	0.1	1	1	0.2
	f_a^{crav}	1	0.2	1	1	0.2	1	0.1	1	1	0.2
Performance, $t \geq$ TQD	t_r^{cpd}	2	2	2	2	2	2	2	1	1	1
	e_{50}^{cpd}	0.00	0.00	0.00	0.00	0.00	13.52	13.52	0.00	0.00	0.00
	e_{50}^{crav}	4.70	2.38	16.32	16.86	12.67	15.05	15.05	2.21	2.45	0.00
	Total Cigs	16.96	16.96	40.23	40.23	39.12	402.44	402.15	20.33	20.33	20.33
	Days CPD=0	34	34	32	32	34	1	1	35	35	35
	$max(cpd)$	10.60	10.60	30.00	30.00	30.00	25.63	25.63	20.33	20.33	20.33
	$max(crav)$	28.68	28.68	31.05	31.05	17.99	32.58	32.58	24.38	24.38	24.38
	$max(WIP_T)$	22.00	22.00	18.00	17.00	21.00	22.00	22.00	6.00	3.00	22.00
	J_e^{cpd}	152.82	152.82	997.26	997.26	983.14	5286.42	5274.37	413.15	413.15	413.15
	J_e^{crav}	2242.77	2312.65	7744.45	7951.02	3889.60	6345.40	6327.91	1660.81	1908.47	1182.94
	$J_e^{wip_T}$	8865.00	10546.00	4508.00	4068.00	5678.00	15529.00	15550.00	87.00	63.00	1836.00
	$var(cpd)$	4.14	4.14	27.21	27.21	26.88	22.50	22.34	11.48	11.48	11.48
	$var(crav)$	33.20	38.70	59.73	54.85	10.20	47.35	47.31	34.48	29.86	31.55
	$var(WIP_T)$	26.12	34.33	33.79	30.84	23.56	35.68	35.14	0.88	0.33	21.34
	$J_I^{u_c}$	5	5	3	3	4	5	5	3	2	4
	$J_I^{u_b}$	118	100	115	115	109	127	127	34	34	34
	$J_I^{u_l}$	7192	8916	3293	2920	4462	13199	13216	29	12	1423
	$var(u_c)$	0.12	0.12	0.08	0.08	0.10	0.12	0.12	0.08	0.05	0.10
	$var(u_b)$	0.32	0.37	0.52	0.52	0.35	0.28	0.28	0.05	0.05	0.05
	$var(u_l)$	18.03	23.31	26.88	24.71	21.17	21.91	21.33	0.76	0.29	19.05

Subject 4 essentially quits smoking on their own and features a significant reduction in *Craving*. Figures 4.18 through 4.21 depict the dosing demands and outcome responses for a subset of tuning combinations considered. The open-loop step responses of these alternate subject models are depicted as the dashed grey responses in their respective intervention simulations. Note, there is no mismatch in the dose-response models in these simulations (other than that due to the causality issue in the dose-response models for the lozenge, see Fig. 3.4). The nominal *Quit*-path models remain

to be those describing the self-regulatory dynamics for the representative patient per equations 2.35 and 2.36, although the K_d parameter in the nominal models (see equation 2.31) is set to each subject's respective baseline CPD level. The alternate subjects' set points and controlled variable lower limit constraints are adjusted to the respective average pre-TQD baseline levels appropriately. Table 4.13 contains the performance metrics for a selection of simulations considered in these analyses.

Subject 1

Subject 1's baseline CPD level is slightly larger than that of the representative patient; Subject 1's baseline $Craving$ is approximately 50% larger than that of the representative patient. Fig. 4.18a depicts the dosing and responses where $f_a^{cpd} = f_a^{crav} = 1$. Aggressive dosing is seen early on in each manipulated variable, which quickly brings down the peak and overall level of $Craving$. The intervention reduces u_c and u_l dosing within days of TQD, as maximum dosing levels are not needed to support tracking of the CPD target. However, high lozenge dosing is implemented and sustained as of day 34 in an attempt to mitigate the high $Craving$ levels that would otherwise be observed.

As Subject 1 is able to track the CPD target with less treatment than the nominal case, Fig. 4.18a reflects relative poor balance between the general goals of avoiding over-dosing and mitigating $Craving$ as a risk factor. However, detuning has a limited effect. Adjusting either the α_r^{cpd} or f_a^{cpd} parameter has little effect as relatively little effort is required to quickly reach the CPD target, regardless of whether the controller is tuned to pursue this target aggressively or otherwise. Detuning via the α_r^{crav} or f_a^{crav} parameters only balances the $Craving$ and dosing trade-off for a period of time. This is apparent with Fig. 4.18b. As the 3DoF tuning parameters slow the rate at which a target is tracked and mismatch rejected, detuning affects dosing for only a

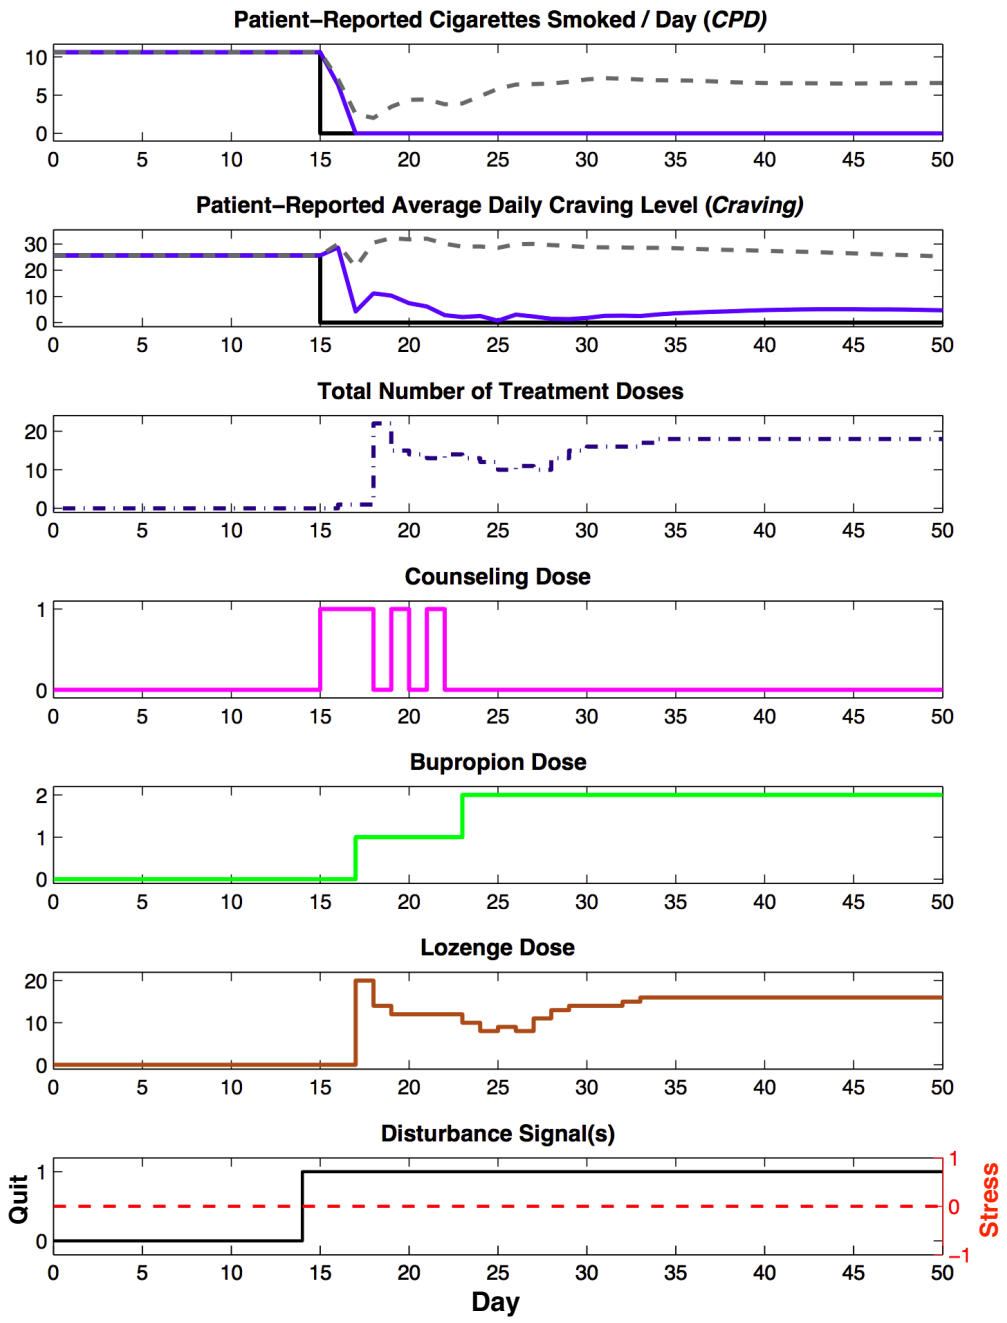

(a) Subject 1: Dosing and Responses (Solid Blue) Where $f_a^{cpd} = f_a^{crav} = 1$. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

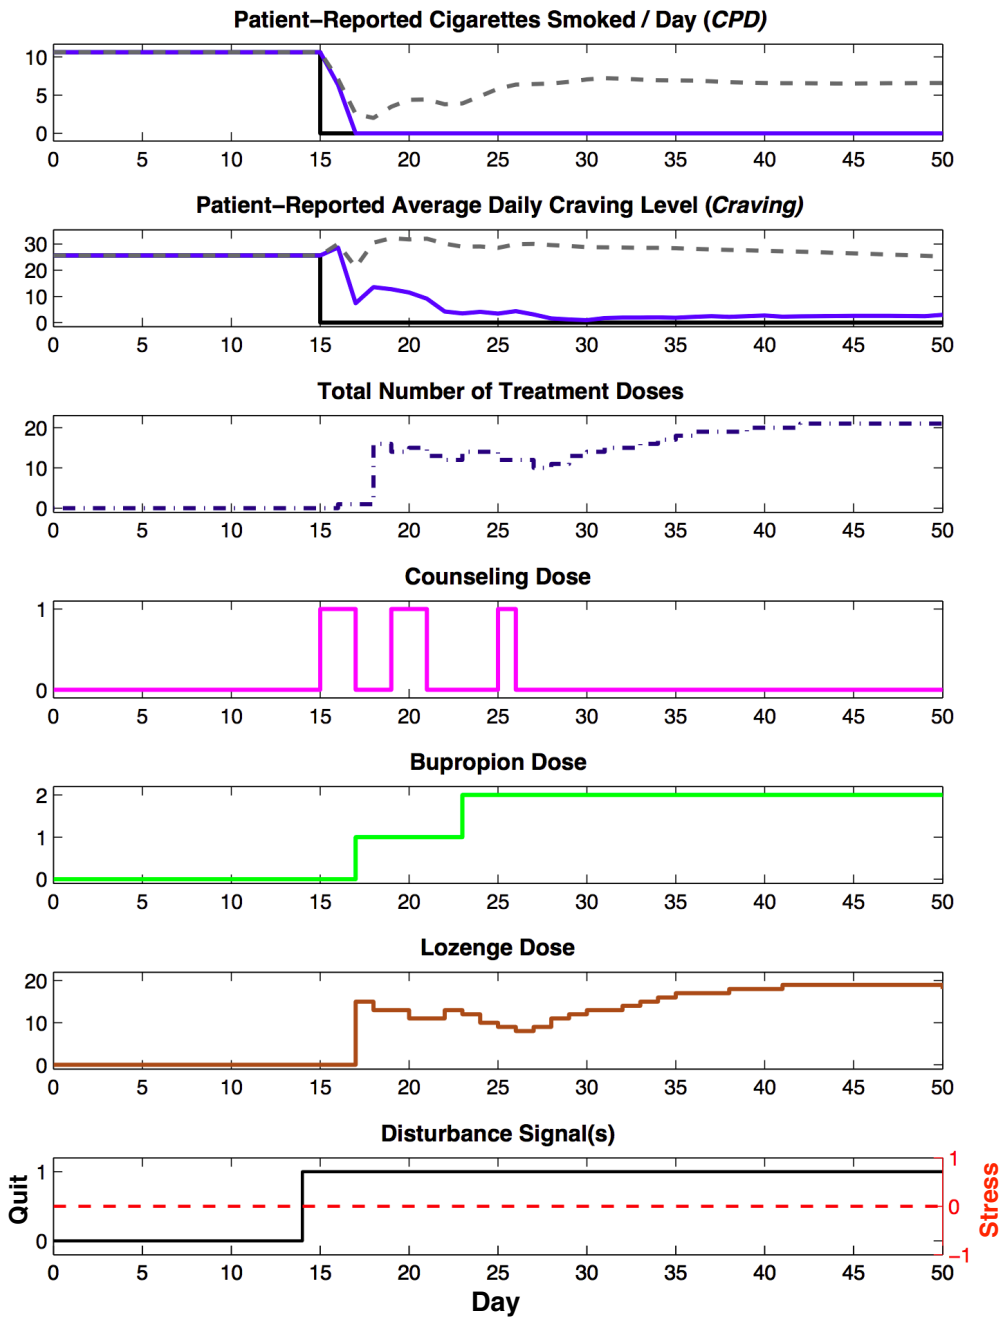

(b) Subject 1: Dosing and Responses (Solid Blue) where $f_a^{cpd} = 1, f_a^{crav} = 0.2$. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

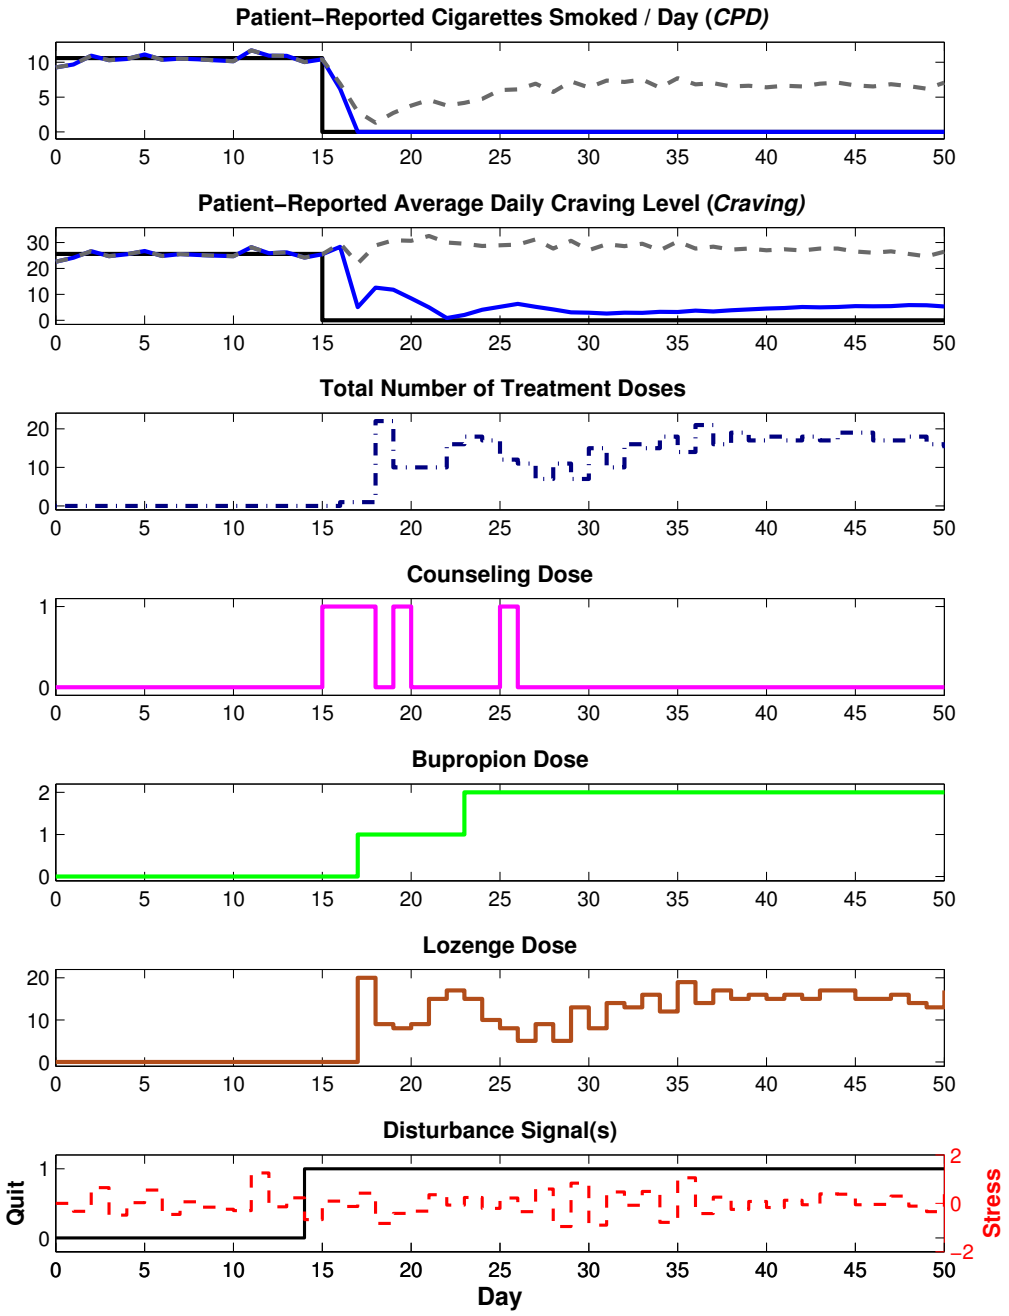

(c) Subject 1: Dosing and Responses (Solid Blue) where $f_a^{cpd} = 1, f_a^{crav} = 0.2$ and a Stochastic *Stress* Disturbance is Present. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

Figure 4.18: Subject 1: Dosing and Responses (Solid Blue) for Select Observer Gain Tuning Levels. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

period of time in the first few weeks of TQD, and high dosing levels are assigned after longer periods of time to decrease the *Craving* offset that would otherwise remain large. Fig. 4.18c illustrates that similar lozenge levels are assigned after several days when a stochastic *Stress* disturbance is present. Even despite the presence of this disturbance, the intervention is able to support a successful quit attempt in terms of CPD, but does lead to slightly higher *Craving* offset.

Altogether, Fig. 4.18 indicates that intervention performance is relatively robust to this sort of mismatch in terms of the CPD and *Craving* outcomes, but at the expense of the WIP_T target. Specifically, it appears that the extent of aggressive dosing is sensitive to the baseline *Craving* level.

Subject 2

Subject 2's baseline CPD level is more than three times that of the representative subject. Although the *Craving* baseline levels are similar, the open-loop response of *Craving* to *Quit* consists of a sustained increase in *Craving* of approximately 50%.

Fig. 4.19a depicts the dosing and responses where $f_a^{cpd} = f_a^{crav} = 1$. The intervention is hesitant just after TQD, as seen in most nominal performance simulations, which serves to avoid large amounts of relaxation of the lower bound on CPD early in the quit attempt. An aggressive lozenge regimen begins on day 20 in order to reduce Subject 2's very high *Craving* levels. Interestingly, the intervention ultimately decreases lozenge dosing starting approximately 18 days after TQD, which corresponds to a gradual increase in *Craving*, although the CPD target is still tracked. This unintuitive control action is the result of a number of factors. Notably, this can be partially attributed to a trade off in the objective function between dosing and constraint relaxation, as large dosing would move CPD far beyond the bound according to the nominal models. Also, CPD features a larger baseline level than

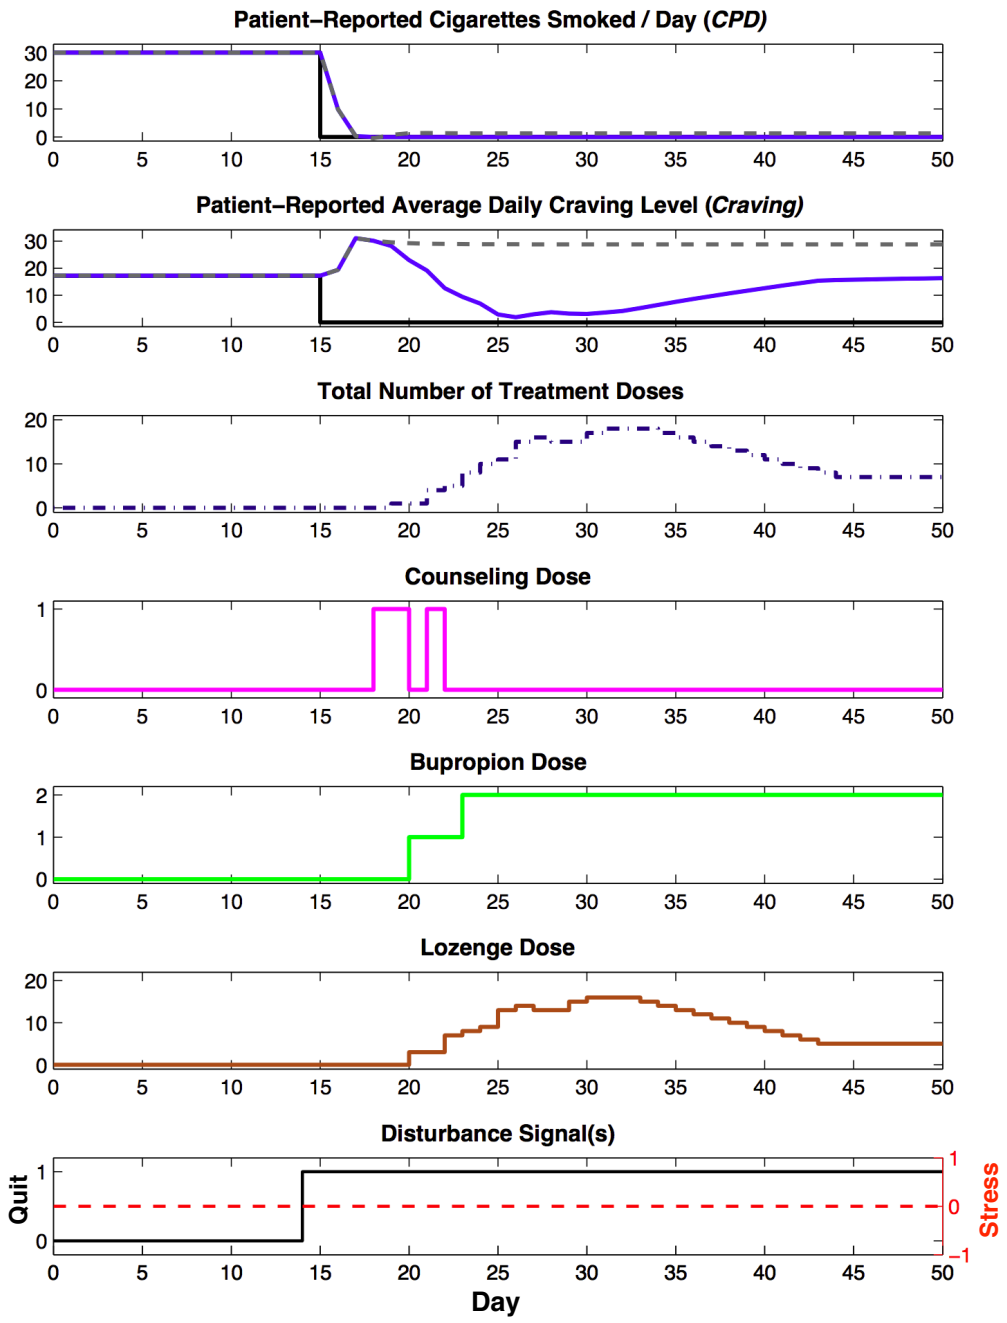

(a) Subject 2: Dosing and Responses (Solid Blue) where $f_a^{cpd} = f_a^{crav} = 1$. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

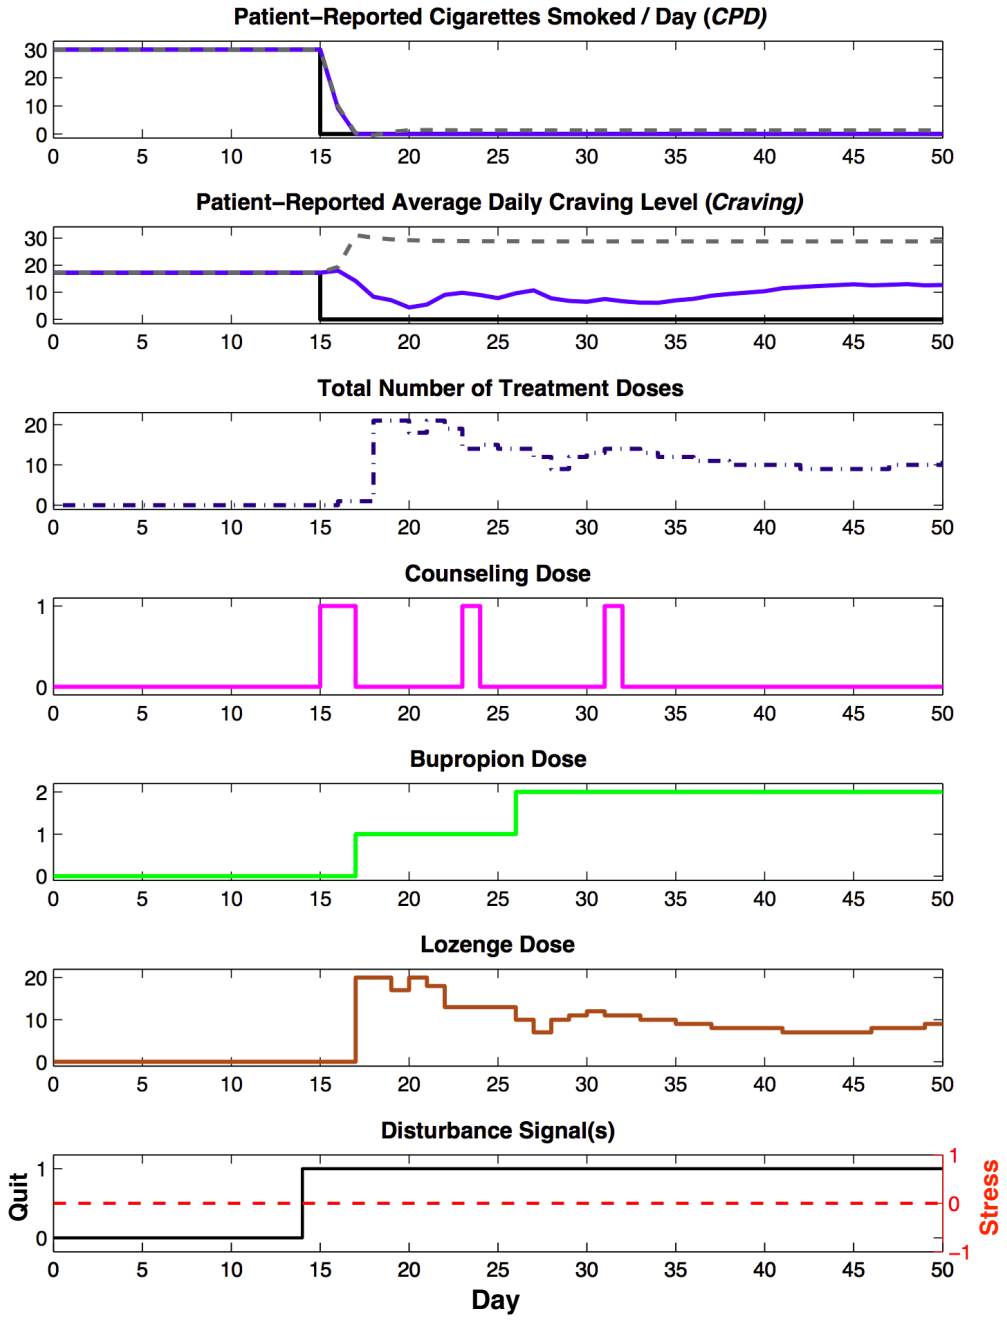

(b) Subject 2: Dosing and Responses (Solid Blue) where $f_a^{cpd} = 1$ and $f_a^{crav} = 0.2$. Open-Loop Responses of CPD and $Craving$ to $Quit$ are also Depicted (Dashed Grey).

Figure 4.19: Subject 2: Dosing and Responses (Solid Blue) for Select Observer Gain Tuning Levels. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

Craving, but the *CPD* and *Craving* targets are penalized equally in the objective function. This means that successful tracking of the *CPD* target leads to lower minimum J values than would perfect tracking of the *Craving* target. This assertion is supported by the contrasting dosing assignments after day 35 for Subjects 1 and 2 when $f_a^{cpd} = f_a^{crav} = 1$ (see Fig. 4.18a and Fig. 4.19a, respectively). The baseline *Craving* level is larger than the baseline *CPD* level for Subject 1, meaning tracking the *Craving* target is "harder" than tracking the *CPD* target, hence the consistently aggressive dosing after day 35 for Subject 1 as tracking the *Craving* target perfectly will reduce J values more than perfect tracking of Subject 1's *CPD* target.

Independently decreasing the f_a^{cpd} value leads to more consistently aggressive lozenge dosing for longer periods of time as of day 20 (not shown), a result of the increased priority of aggressively tracking the *Craving* target. Detuning via independent decreases in f_a^{crav} actually improves *Craving* performance early on. As dosing decisions focus on tracking of the *CPD* target when $f_a^{cpd} = 1$ and $f_a^{crav} < 1$, the intervention aggressively doses early in order to suppress the nearly 30 cigarette resumption that the nominal model predicts would otherwise occur; this has a simultaneous effect of reducing the peak *Craving* levels. Both aspects of these independent detuning scenarios are observed in Fig. 4.19b, where $f_a^{cpd} = f_a^{crav} = 0.2$. As before, though, detuning effects are most distinct for the first few weeks immediately following TQD only, reflected in the similar *Craving* offsets on day 50 for Fig. 4.19a and Fig. 4.19b.

Subject 3

Subject 3's baseline *CPD* and *Craving* levels are both significantly higher than those for the representative subject. This, and the fact that Subject 3's open-loop *Quit*-response models indicate only one day of decreased *CPD* levels during the quit attempt and no decrease in *Craving* levels, leads to aggressive dosing in all components

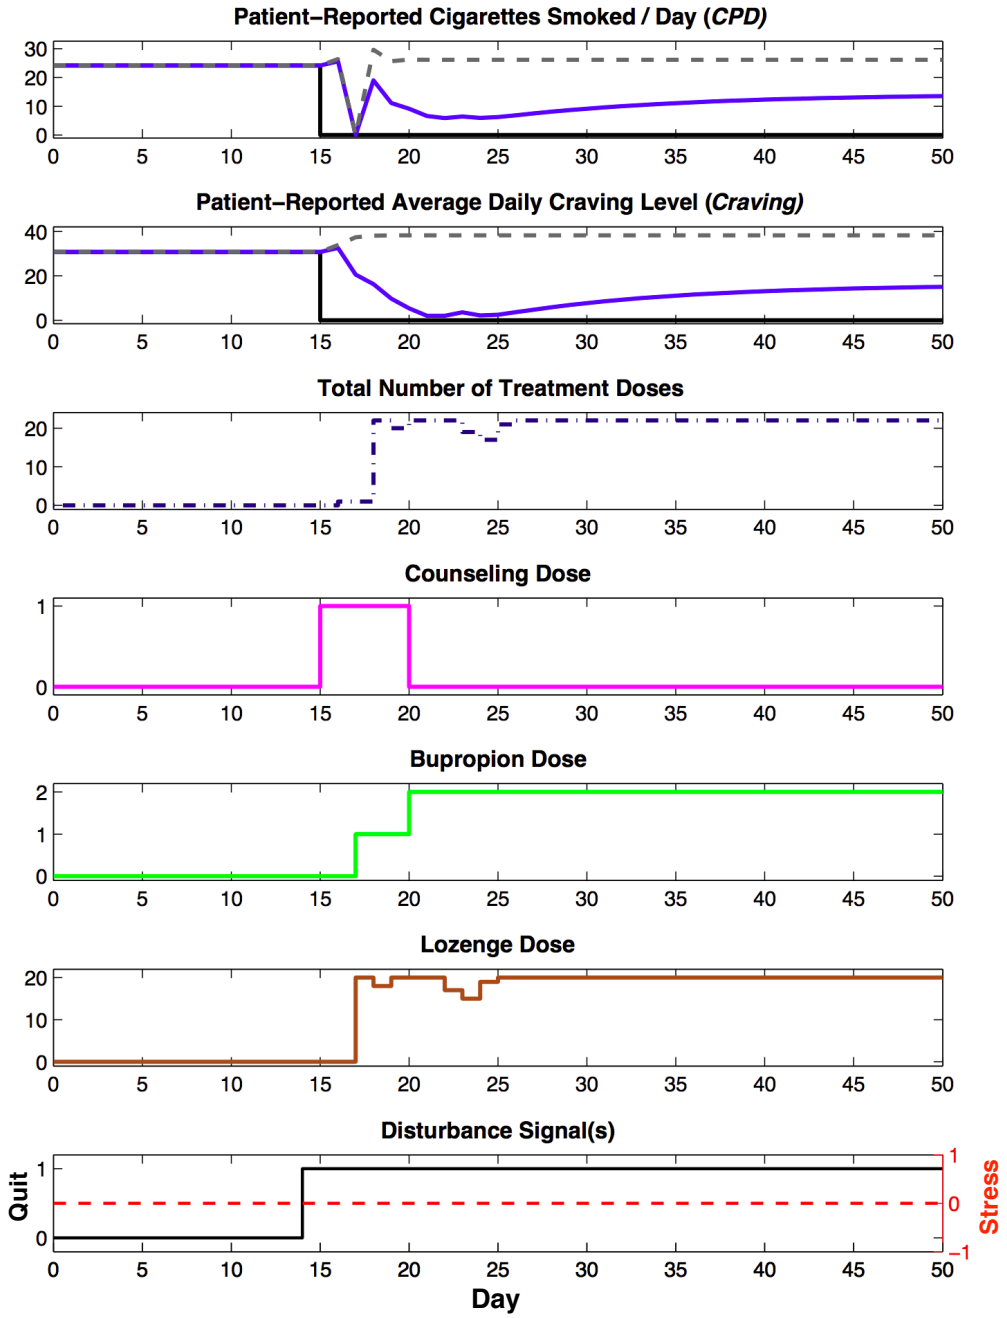

Figure 4.20: Subject 3: Dosing and Responses (Solid Blue) where $f_a^{cpd} = f_a^{crav} = 1$. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

early on. Depicted in Fig. 4.20, this aggressive dosing is generally sustained for the duration of the quit attempt. Such dosing does move both CPD and $Craving$ to approximately 50% of the baseline levels as of day 17, but the patient still features nontrivial smoking for most of the quit attempt, and $Craving$ remains at approximately 18 points as of day 50. Detuning via any of the tuning parameters may suppress aggressive dosing early on, but at the expense of the CPD and $Craving$ outcomes, and ultimately still leads to an unsuccessful quit attempt.

Subject 4

While Subject 3 represents a relatively extreme level of mismatch, Subject 4 represents an alternative mismatch extreme in that the patient can successfully reduce smoking and $Craving$ on their own over time. Subject 4's baseline CPD and $Craving$ levels are both larger than those for the representative subject. As seen in Fig. 4.21, the intervention adapts appropriately to the patient's early and consistently successful attempt to quit smoking: counseling is only briefly assigned just after TQD, only one bupropion dose is assigned, and few lozenges are assigned. This minor dosing serves to bring $Craving$ to trivial levels by day 21, which still corresponds to successful tracking of the CPD target. Altogether, Fig. 4.21 highlights the appealing resource use aspect of this approach, as the dosing depicted in the plots is less than that generally assigned by current clinical recommendations, yet successful outcomes are still achieved and maintained. Because of the success of the patient's open-loop quit attempt, the intervention does not necessarily benefit from detuning. Even extreme detuning pursued by significantly increasing α_r^{cpd} and α_r^{crav} corresponds to little change in dosing demands compared to the $\alpha_r^{cpd} = \alpha_r^{crav} = 0$ case; the level of $Craving$ for the detuned scenario where $\alpha_r^* > 0$ is actually slightly higher beginning around day 17 to day compared to the $Craving$ observed for the $\alpha_r^{cpd} = \alpha_r^{crav} = 0$ case. Adjusting f_a^{cpd}

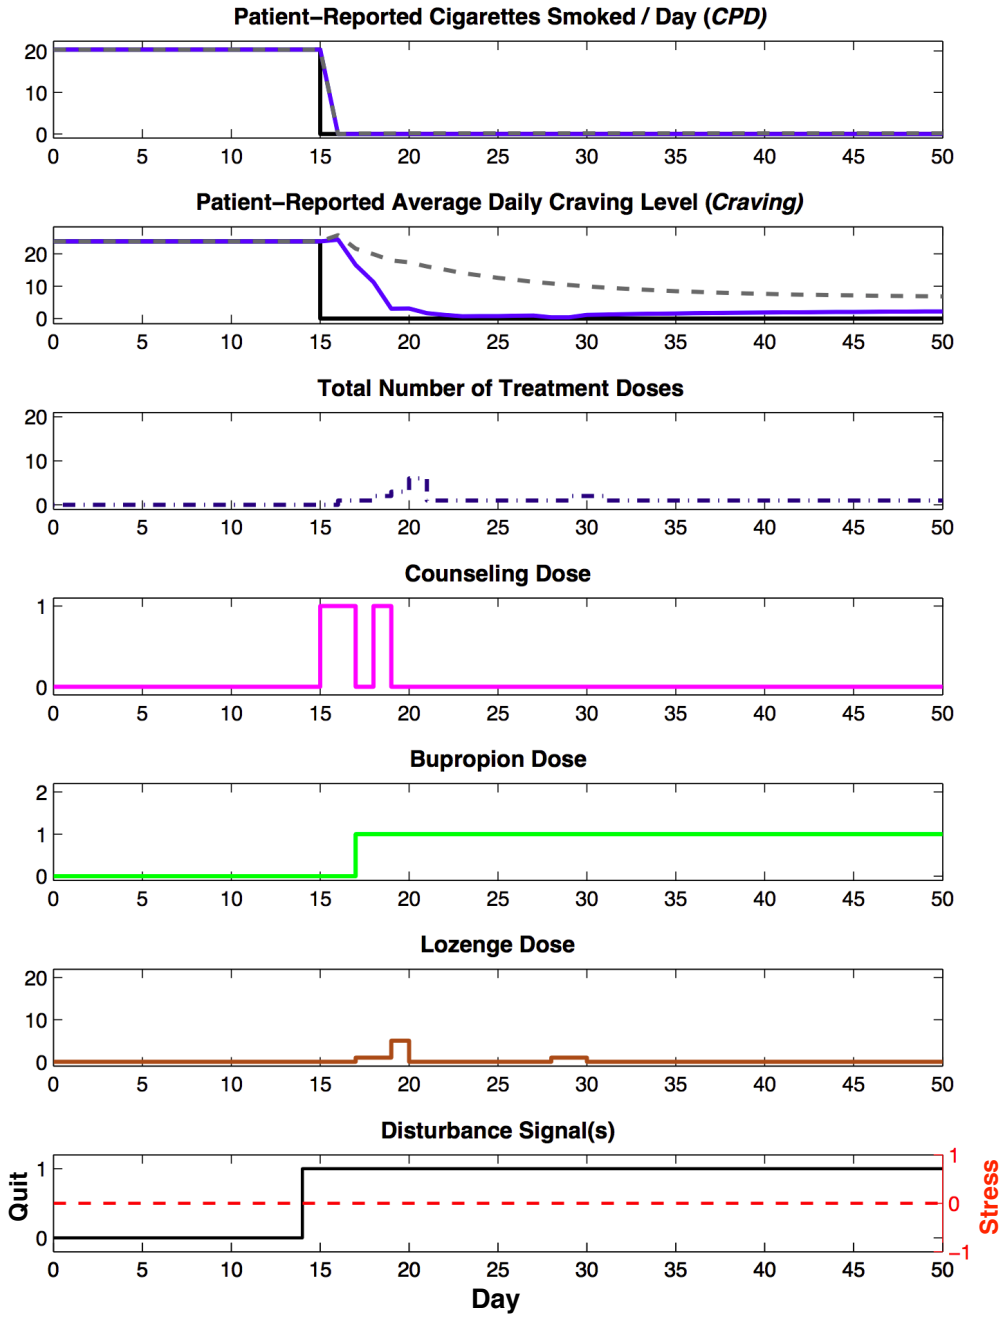

Figure 4.21: Subject 4: Dosing and Responses (Solid Blue) where $f_a^{cpd} = f_a^{crav} = 1$. Open-Loop Responses of CPD and $Craving$ to $Quit$ are Also Depicted (Dashed Grey).

or f_a^{crav} independently or simultaneously actually leads to more aggressive dosing overall, particularly within the first 10 days of the quit attempt. As an example, independently detuning such that $f_a^{cpd} = 0$ and $f_a^{crav} = 0.2$ leads the intervention to assign the maximum u_b dose as of day 20 as well as some lozenge. This occurs as $f_a^{crav} < 1$ leads the observer to trust the nominal *Craving* models more than the *Craving* measurement, and the nominal *Craving* models suggest *Craving* will remain high after TQD in the absence of relatively aggressive dosing.

4.5.4 Summary

The scenarios presented in this section indicate that the performance of this intervention approach can be robust in the presence of a variety of noise and uncertainty conditions. Similarly, the effective adaptation that balances tracking of intervention targets with resource use considerations suggest that the character of the representative patient's cessation dynamics offer good initial nominal models.

While aggressive dosing often leads to tracking of the intervention targets, the 3DoF tuning capabilities can offer clinicians a flexible way to obtain dosing schedules suited to individual patient circumstances. The simulated results appear to indicate that patients for whom adherence is an issue may benefit most from a detuned scenario, as detuning generally leads to lower peak lozenge demands and more subtle lozenge dosing adjustments in particular. Specifically, decreasing the tuning parameters in the observer gain matrix offer improved dosing and potentially improved outcomes for tuning conditions including the following: when significant noise is present in *Craving* measurements, where decreasing f_a^{cpd} specifically leads to improved dosing and even CPD dynamics (see Fig. 4.15d and Fig. 4.15a); when aggressive bupropion dosing is more favorable than high lozenge levels; and when there is a benefit to aggressive dosing early in the intervention over later in the intervention or vice-versa,

as detuning in the presence of mismatch can manage a trade-off between aggressive tracking of the CPD target early in the quit attempt, versus aggressive tracking of the $Craving$ target as time goes on. Detuning appears to negatively influence dosing assignments when the patient receiving the intervention features open-loop cessation dynamics that independently lower CPD and $Craving$ levels.

That said, robust performance appears to be most sensitive to uncertainty in the effectiveness of bupropion (see Section 4.5.2) and the steady-state levels of CPD and $Craving$ in the open-loop. Diverging open-loop steady-state levels can bias the intervention toward aggressive dosing that tracks one target over another weeks after TQD, based on how much tracking one set point contributes to reduction of the objective function values at the decision points. In other words, the controller formulation evaluated in this section is particularly sensitive to the relative steady-state levels of the $Quit$-path processes in the open-loop. High steady-state $Craving$ levels in the open-loop relative to that patient's steady-state CPD levels in the open loop, and to that of the nominal patient's $Craving$ levels appear to have a significant effect; this scenario tends to assign large amounts of all treatment components even weeks after TQD, regardless of success in CPD target tracking. As detuning only slows the speed at which the targets are tracked or mismatch rejected, doing so in this case still leads to high dosing at the end of the time period examined. This suggests that future research may benefit from development of a formulaic method to determine a more advantageous ratio of Q_{cpd} to Q_{crav} as a function of the relative baseline levels. However, this issue speaks more to the need for accurate dose-response models as Section 4.5.3 does not employ dose-response models other than those describing the representative, yet treatment effects are likely a function of levels of addiction and other factors; this concept of dosing effects being a function of addiction level is even addressed in current clinical practice as 4 mg nicotine lozenges—as opposed to 2 mg

lozenges—are assigned to heavier smokers (Rennard *et al.*, 2014).

Chapter 5

CONCLUSIONS & FUTURE WORK

5.1 Summary

This dissertation has explored the usefulness of applying process modeling and control concepts from engineering into the study of cigarette smoking behavior change and the development of an adaptive smoking intervention. Toward this goal, two distinct but related aspects of this research were presented: (1) the use of dynamical systems modeling and system identification methods to describe behavior change during a smoking cessation intervention (Chapter 2), and (2) the formulation (Chapter 3) and evaluation (Chapter 4) of a decision framework based in Hybrid Model Predictive Control (HMPC) that adapts treatment components to meet the changing needs of a patient.

5.1.1 Modeling Smoking Cessation Behavior Change

Two hypothesized mechanisms of change that are central to the study of cause and effect in behavioral and social science settings were cast as dynamical systems in this research. Specifically, dynamical systems models were developed to describe behavior change according to statistical mediation and self-regulatory mechanisms. As was illustrated in Chapter 2, these models examine relationships between time-varying behaviors as input-output systems; system identification methods offer a means to utilize behavioral ILD to describe these dynamic behaviors with low order differential equation models. The major contributions of this modeling work include:

- Demonstration that an engineering modeling approach can be used to effectively

understand dynamic behavioral phenomena: The work described in Chapter 2 illustrates that linear time-invariant models in continuous-time ODE form and prediction-error identification methods offer a means for obtaining clinically-meaningful models of behavior change. This modeling paradigm can provide significant insight into tobacco-use behaviors in particular. Models were estimated using both treatment group average and single subject data sets. This suggests this modeling methodology can be used to investigate broad research questions within behavioral science.

- Development of a dynamical systems representation of statistical mediation: The dynamic mediation models presented here stem from a connection to production-inventory systems and are not necessarily limited to the study of cigarette smoking. These models represent each path in a mediated behavioral system as an independent input-output process, which together describe mediated behavior change as a parallel-cascade system.

- Estimation and interpretation of dynamic mediation models that describe cessation-induced changes in CPD as a $Craving$-mediated process: Using ILD collected in the McCarthy *et al.* (2008b) clinical trial, treatment group average continuous-time ODEs were estimated. Estimation of each set of expressions consists of estimated low-order transfer functions and high goodness-of-fit values. These models attribute the large, immediate drop in CPD on TQD to the non-mediated path and the relatively small and slow resumption of smoking after TQD to changes in $Craving$. Comparing the parameter estimates across the groups suggests counseling and bupropion independently and in combination affect the cessation process, generally influencing the net change in $Craving$, the degree to which CPD decreases on TQD, and the magnitude and speed at which CPD

increases after TQD. Ultimately, these models suggest an interrelationship between *Craving* and *CPD* cannot be fully described by one mediation model structure.

- Reverse-engineering of models describing a self-regulating psychological process: Through secondary analysis of ILD, empirical closed-loop models were estimated and validated to represent self-regulation within behavior change. Models of this nature have historically been difficult to precisely quantify with traditional behavioral science analytical methods, and this mechanism was previously described largely in conceptual terms in behavioral science and substance use literature (Carver and Scheier, 1998; Solomon, 1977; Solomon and Corbit, 1974; Velicer *et al.*, 1992).

- Development of models describing smoking cessation behavior change as a self-regulatory process: A specific self-regulation system was proposed and modeled in this work. This feedback process represents *Craving* as the controlled variable, *CPD* as the input variable to the process where *Craving* is the output. *CPD* itself is the result of an excitatory *Quit* disturbance as well as the output of a biological or psychological *Craving* self-regulator.

- Estimation and interpretation of models quantifying self-regulation within day-to-day changes in *Craving* and *CPD* during a quit attempt: Continuous-time ODE models were estimated using data from the McCarthy *et al.* (2008b) study through a two step closed-loop identification problem. They suggest that on average, the *Craving* set point is the average *Craving* level pre-TQD and that the self-regulator acts as a proportional-with-filter controller.

Altogether, this dissertation makes the case for the usefulness and relevance of dynamical systems modeling and system identification techniques for better under-

standing dynamic health behaviors generally, and specifically for providing clinically-meaningful insight into the nature of cigarette smoking behavior change.

5.1.2 Adaptive Intervention Design & Evaluation

Adaptive smoking cessation interventions seek to explicitly address the chronic, relapsing nature of tobacco dependence by adjusting treatment component dosages to the changing needs of a specific patient. A control systems engineering paradigm for designing such interventions offers a means for systematic treatment optimization and personalization. In Chapter 3, an adaptive intervention algorithm employing an HMPC framework was developed. This intervention was evaluated extensively through simulation, as described in Chapter 4. Ultimately, this work offers a proof-of-concept for the potential clinical application of an engineering approach to the design of cigarette smoking treatments. The major contributions of this intervention design work include:

- Identification of the major requirements of a clinically-relevant adaptive cessation intervention and establishment of a connection to control systems engineering ideas: Connections between specific clinical considerations and control systems engineering ideas were outlined; these clinical circumstances and demands manifest in a control systems framework as controlled variables, controller structure, constraints, manipulated variable constraints, and more.

- Development of a mathematical representation of the nominal model: A set of open-loop *Quit*-response, *Stress*-response, and dose-response models were obtained to represent a hypothetical patient unable to quit smoking on their own. The discrete-valued nature of treatment component dosages were represented as linear inequality constraints via continuous and discrete auxiliary variables

incorporated into the MLD model framework. Discrete-time, state-space representations of these models were then presented.

- Formulation of an HMPC-based adaptive cessation intervention: The general requirements, goals, and constraints of a clinically-relevant control-based interventions were translated into mathematical form and incorporated into the HMPC framework. The result is an intervention algorithm that determines daily adjustments to counseling, bupropion, and nicotine replacement lozenge dosages in order to move controlled variables to their target levels while rejecting the influence of a measured *Stress* disturbance and accounting for an array of hard constraints. 3DoF tuning capabilities were also incorporated.

- Demonstration that good therapeutic performance can be obtained in the nominal case: Diverse simulations indicate that an HMPC-based intervention can promote successful cessation in an individual unable to do so on their own. The objective function penalty weights and 3DoF tuning parameters offer a means for obtaining dosing profiles and post-TQD CPD and *Craving* responses of diverse character.

- Illustration that the therapeutic performance of an HMPC-based smoking cessation intervention can be robust to measurement noise, parametric uncertainty, and significant plant-model mismatch: Among other findings, simulations that assessed robust performance suggest that tuning via the observer gain offers significant improvement in terms of dosing demands and outcomes in the presence of measurement noise. Additionally, dosing demands appear to be largely sensitive to a patient's baseline *Craving* level, particularly its value relative to CPD baseline levels and to the nominal patient's baseline levels.

5.1.3 Future Work

A number of opportunities and lines of research exist to move the ideas and analyses presented in this dissertation toward clinical practice.

Improving Model & Controller Complexity

While not an exhaustive list of opportunities, a number of natural extensions of the work presented are described below. Significant progress on this work can likely build on existing work in a relatively straightforward manner and in the absence of novel experimentation.

Given the nonlinear nature of physiological phenomena and the diverse manner in which inter-individual variability can enter determinants of human behaviors, incorporating additional complexity into the models presented here may result in expressions representative of dynamics observed over a wider range of circumstances or time scales. For example, slowly time-varying or static variables, such as demographic variables, likely play a nontrivial role within an individual's cessation process. The systematic effect of such factors could be incorporated into future self-regulation models. As an example, a linear parameter-varying form of the closed-loop models presented in Chapter 2 could be developed in which a scheduling parameter may represent an individual's socioeconomic status, for example (Novara *et al.*, 2011).

Furthermore, the dynamic psychological component of the cessation process is intuitively nonlinear as there is physical and logical lower bound on both CPD and $Craving$. In other words, the fact that CPD and $Craving$ cannot realistically assume negative values indicates that the open-loop behavior change process actually consists of a saturated system (Widanage *et al.*, 2004). While such saturation was incorporated into simulations of the patient-reported responses, this factor was not

incorporated into the nominal models. Future efforts may benefit from incorporating this nonlinearity into the open-loop models that act as the basis for intervention design.

Formulaically, the intervention developed here can be altered in a straightforward manner to include additional or alternate treatment components, controlled variables, and disturbances. For example, self-efficacy and negative affect are thought to play significant roles in the cessation process (Gwaltney *et al.*, 2009; Piper and Curtin, 2006; Shiffman and Waters, 2004) and could be incorporated as additional controlled variables. Similarly, an alternate set of treatment components could be employed. For example, nicotine replacement patches are the most commonly used treatment product (Reynolds, 2011). Future work may want to consider combination therapy consisting of a nicotine replacement patch in lieu of bupropion as the former is more widely available. Reformulating the intervention to alter the combination of manipulated and/or controlled variables would be relatively straightforward. However, without additional experimentation, the open-loop models around which an HMPC-based intervention is designed will be limited by the quality of the dynamical systems models.

Incorporating Within-Day Decision-Making

Design of the intervention in Chapter 3 with daily measurements and adaptation is motivated by a number of factors. First, the models in Chapter 2 indicate relevant behavioral phenomena occur on the daily time scale. Additionally, dosages of many existing treatment components can only realistically be adjusted from day-to-day (or even more infrequently), and their effects are more appropriately described by a daily time scale; counseling, bupropion, and nicotine replacement patches fall into this category. Furthermore, daily measurement and adaptation requirements should be

manageable from a logistical perspective.

However, a controller involving within-day measurement and treatment adaptation may improve how effective lozenge assignments are in terms of preventing specific instances of smoking. An intervention algorithm of this sort may also be able to mitigate issues of plant-model mismatch for the u_l and $Stress$ effects, which were introduced to the controller formulation in Chapter 3 in order to maintain a causal controller (see Section 3.3.2). That said, simply extending the HMPC formulation described in Chapter 3 such that sampling and control action decisions occur on a minutely or hourly basis would result in an unnecessarily complex intervention algorithm with tractability issues. Specifically, a single overarching control algorithm that is meant to determine counseling, bupropion, and lozenge control action multiple times within a day would mean the prediction horizon value (p) employed in Chapter 4 translates to a p in this hypothetical scenario that is an order of magnitude larger. These large prediction and move horizons would also include spans of time where measurement and adaptation is not feasible (e.g., while a patient is sleeping). Furthermore, this more complex controller would also require switching time constraints such that counseling and bupropion control action could not be implemented more than once daily (or even less in the case of bupropion). Instead, a hierarchical control scheme can offer a means to manage the multi-timescale nature of the cessation process, dosing frequency, and speed of treatment effects in a conceptually and computationally simpler manner (Barcelli, 2012; Brdys *et al.*, 2008; Scattolini and Colaneri, 2007; Van Henten and Bontsema, 2009). A survey of the various forms a hierarchical control scheme can take is outside the scope of this chapter; Barcelli (2012) and Scattolini and Colaneri (2007) offer good overviews of this class of control framework.

In the context of an adaptive smoking intervention, a decision framework that

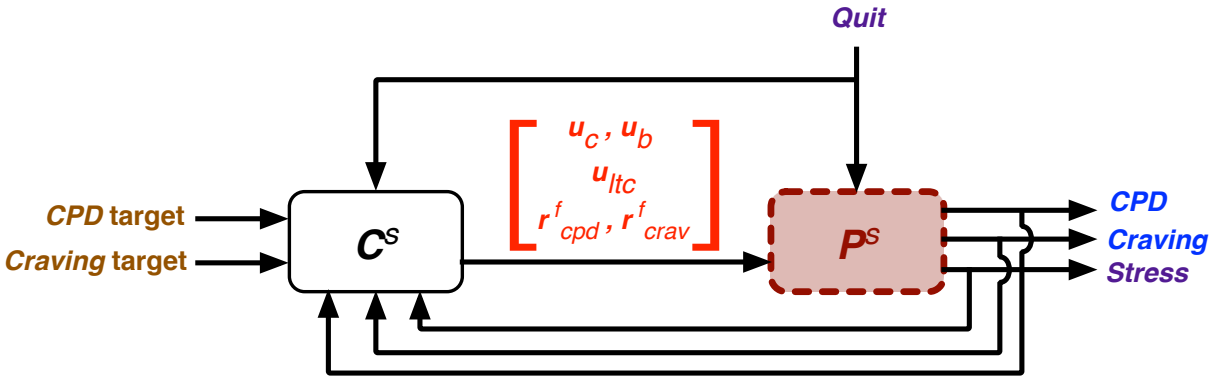

(a) Block Diagram for the Upper, Slow Level of the Hierarchical Control Scheme. P^s Encompasses the Fast Level of the Control Scheme Depicted in Fig. 5.1b.

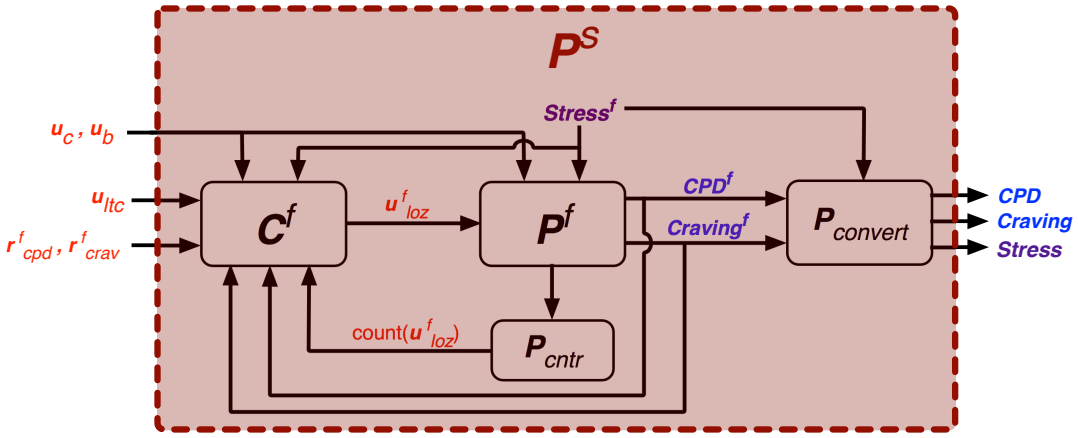

(b) Block Diagram for the Within-Day, Fast Level of the Hierarchical Control Scheme. This Subsystem Constitutes the Process Block in the Block Diagram of the Slow Level of the Control Scheme, i.e., P^s in Fig. 5.1a.

Figure 5.1: Block diagram of a Potential Hierarchical Control Scheme for an Adaptive Intervention that Features Combined Within-Day and Between-Day Dosing.

features within-day decision making may take the form of the block diagrams in Fig. 5.1. Generally, this framework considers a "slow" level that operates on a daily time scale and a "fast" level that operates on a within-day time scale (e.g., minutes, half hours, multi-hour, etc.). The slow controller defines dosing for the treatment components whose dosages can only be delivered on a daily basis and/or whose effects are observed more on a multi-day time scale. It also determines the total number of lozenges to be taken during the following day, and this value acts as a terminal constraint on the fast time scale.

More specifically, P^s in Fig. 5.1a represents the behavior change system from the higher level (the "slow" level, indicated with the s superscript). On this slow level, $Quit$, u_c, and u_b are the same signals as defined in Chapters 3 and 4. CPD, $Craving$, and $Stress$ represent total daily smoking level, average daily craving level, and average daily stress level, as before, but are no longer patient-reported. Instead, they are constructed from smoking, craving, and stress levels that are reported by the patient multiple times within a given day; this is why $Stress$ is an output of P^s in Fig. 5.1a, but it will not act as a controlled variable. The slow controller, C^s, is provided with the intervention targets and values of $Quit$, $Craving$, and $Stress$ each day to determine adjustments to u_c, u_b, the number of lozenges to be taken in the day (u_{ltc}, which will act as a terminal constraint in the within-day controller), and the CPD and $Craving$ set points (r^f_{cpd} and r^f_{crav}, respectively). (These set points will likely equal the CPD and $Craving$ targets defined for the slow level—i.e., = baseline levels, $t <$ TQD, and $= 0$, $t \geq$ TQD—but this may not necessarily be true if the intervention target is not full cessation or if detuning tracking of the targets is desired.)

On the fast level, depicted in Fig. 5.1b, patient-reported smoking, craving, and stress levels are obtained at some time interval, e.g., every two hours; CPD^f, $Craving^f$,

and $Stress^f$ are measurements of the total number of cigarettes smoked since the last report, average craving level since the last report, and average stress level since the last report, respectively. On this control level, the manipulated variable is u^f_{loz}, i.e., the number of lozenges to be taken consecutively between now and next sampling instance. CPD^f, $Craving^f$, and the total number of lozenges taken since the beginning of the day—count(u^f_{loz})—are the controlled variables. u_c and u_b are defined by C^s, are constant values on this level, and act as feedforward signals; $Stress^f$ is also a feedforward signal. The fast controller C^f determines adjustments to u^f_{loz} such that CPD^f tracks r^f_{cpd}, $Craving^f$ tracks r^f_{crav}, and count(u^f_{loz}) equals u_{ltc} by the end of the day; the values of r^f_{cpd}, r^f_{crav}, and u_{ltc} were defined by C^s. Dose-response and $Stress^f$-response mechanisms of the behavior change phenomena that occurs on this time scale are represented by the plant models, P^f. P_{cntr} is a function counting the total number of lozenges taken since the beginning of the day. $P_{convert}$ translates the patient-reported CPD^f, $Craving^f$, and $Stress^f$ levels into CPD, $Craving$, and $Stress$ constructs (which are used by C^s).

A hierarchical control scheme of this sort could be formulated and evaluated computationally based on the programs and routines employed in this dissertation. However, additional experimentation would be required in order to estimate high-fidelity models of P^s and P^f.

Possible Future Experimentation

This dissertation alluded to a number of questions future experiments may seek to answer. These potential experimental avenues include modeling the natural evolution of craving and smoking levels within a day, between-day dynamics, the self-regulatory system depicted in Fig. 2.7 (regardless of any treatments), validating the novel intervention, and more.

More specifically, future research would significantly benefit from a pilot-type study intended to examine the ideas and findings discussed in this dissertation; in other words, a study conducted prior to a large experiment designed with system identification experimental methods or randomized-controlled structures in mind would be useful. In such a study, modern EMA methods and intensive data collection tools should be employed, as the capabilities of mobile and wearable technologies today are far beyond those commonly available when the McCarthy *et al.* (2008b) study was conducted. Such a pilot study could explore how modern smartphone and mHealth tools can be used to better understand the influence of contextual factors, obtain better single subject models, and more. Furthermore, an initial exploratory study that implements an intervention similar to that described in Chapter 3 could provide insight in to the form a larger randomized, placebo-controlled clinical trial should take, and more generally about the feasibility of an adaptive intervention of this nature.

Ultimately, if the ideas and findings presented in this dissertation are to move toward widespread clinical practice, novel, large-scale clinical studies should have two major goals: (1) estimation and validation of high-fidelity *Quit*-, disturbance-, and dose-response models, and (2) validation of the efficacy of an HMPC-based adaptive intervention relative to current clinical practice.

A novel clinical trial intended to collect ILD better suited for obtaining high-fidelity dose-response dynamical models should be designed with experiment-design methods from system identification in mind. Specifically, input signal design methods should be used to define a "patient-friendly" sequence of treatment dosages that vary over time in a persistently exciting manner (Deshpande *et al.*, 2014; Ljung, 1999). In designing these input signals, one must consider a variety of factors unique to this problem setting. These factors include switching time, amplitude, move size, and integer constraints as dictated by the nature of the treatment components being

examined (Deshpande and Rivera, 2014). For example, the binary nature of u_c implies a counseling input sequence that reflects a random binary signal (RBS) or pseudo-random binary signal (PRBS) may be appropriate. As major swings in lozenge dosing assignments may not be practical clinically, or may threaten subject buy-in, total daily lozenge dosing should likely take the form of a multisine signal. If modeling within-day dynamics is a goal of this future clinical study, the specific time at which a lozenge is assigned within a single day may be determined according to an RBS or PRBS signal, where the total number of lozenges assigned in a given day has been pre-determined according to the multisine signal defining day-to-day total lozenge dosing. Future experiments and interventions may want examine nicotine replacement patches in lieu of bupropion. Compared to bupropion, this form of NRT is more widely used currently (Reynolds, 2011), is more widely available (Rennard *et al.*, 2014), and its dosages can be varied over time more freely (Rennard *et al.*, 2014; Tobacco Use and Dependence Guideline Panel, 2008). While patches are commercially available in 7, 14, and 21 mg of nicotine dosages, not all levels can be assigned to all patients without potentially inducing nicotine toxicity side effects (Rennard *et al.*, 2014; Tobacco Use and Dependence Guideline Panel, 2008). Consequently, a patch protocol may also want to take the form of an RBS or PRBS signal. ILD in this study should be collected via EMA protocol and would likely draw from the same questionnaires employed in the McCarthy *et al.* (2008b) clinical trial. A future study should also be designed to allow cross-validation such that models estimated from a ILD collected from a portion of subjects are validated against ILD collected from the remaining subjects.

Once dynamical systems models with greater predictive ability have been estimated and validated, an HMPC-based adaptive intervention could be formulated with the improved models. A clinical study focusing on the efficacy of such an adaptive intervention could be pursued. Ultimately, this sort of intervention would need to

be examined in a large scale randomized, placebo-controlled study, in which the novel intervention's efficacy is compared to that of current clinical practice, particularly in terms of six month cessation rates.

In conclusion, a number of avenues for future work are apparent that can build from the ideas and findings presented in this dissertation. These future efforts would continue to pursue a better understanding of smoking cessation behavior change and an optimized engineering-based smoking intervention. Ultimately, many additional contributions can be made by future incorporation of dynamical systems and control ideas into the study and treatment of smoking cessation.

REFERENCES

Adsit, R. T., B. M. Fox, T. Tsiolis, C. Ogland, M. Simerson, L. M. Vind, S. M. Bell, A. D. Skora, T. B. Baker and M. C. Fiore, "Using the electronic health record to connect primary care patients to evidence-based telephonic tobacco quiteline services: a closed-loop demonstration project", Translational Behavioral Medicine pp. 1–9 (2014).

Ahmed, S. H. and G. F. Koob, "Transition from moderate to excessive drug intake: Change in hedonic set point", Science **282**, 5387, 298–300 (1998).

Ahn, A. C., M. Tewari, C. Poon and R. S. Phillips, "The limits of reductionism in medicine: Could systems biology offer an alternative?", PLOS Medicine **3**, 6, e208 (2006).

Allen, S. S., T. Bade, D. Hatsukami and B. Center, "Craving, withdrawal, and smoking urges on days immediately prior to smoking relapse", Nicotine and Tobacco Research **10**, 1, 35–45 (2008).

American Cancer Society, "Guide to quitting smoking", URL `http://www.cancer.org/` (2012).

Aveyard, P. and M. Raw, "Improving smoking cessation approaches at the individual level", Tobacco Control **21**, 252–257 (2012).

Baker, T. B., R. Mermelstein, L. M. Collins, M. E. Piper, D. E. Jorenby, S. S. Smith, B. A. Christiansen, T. R. Schlam, J. W. Cook and M. C. Fiore, "New methods for tobacco dependence treatment research", Annals of Behavioral Medicine **41**, 2, 197–207 (2011).

Baker, T. B., M. E. Piper, T. R. Schlam, J. Cook, S. S. Smith, W. Loh and D. Bolt, "Are tobacco dependence and withdrawal related amongst heavy smokers? relevance to conceptualizations of dependence", Journal of Abnormal Psychology **121**, 4, 909–921 (2012).

Barcelli, D., *Decentralized and hierarchical Model Predictive Control or Networked systems*, Ph.D. thesis, Universita Degli Studi Di Siena, Siena, IT (2012).

Baron, R. M. and D. A. Kenny, "The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations", Journal of Personality and Social Psychology **51**, 6, 1173–1182 (1986).

Bekiroglu, K., C. Lagoa, S. A. Murphy and S. T. Lanza, "A robust MPC approach to the design of behavioural treatments", IEEE 52nd Annual Conference on Decision and Control (CDC), 2013 pp. 3505–3510 (2013).

Bemporad, A. and M. Morari, "Control of systems integrating logic, dynamics, and constraints", Automatica **35**, 407–427 (1999).

Benowitz, N. L., "Pharmacology of nicotine: Addiction, smoking-induced disease, and therapeutics", Annual Review of Pharmacology and Toxicology **49**, 57–71 (2009).

Bergen, A. W., H. S. Javitz, R. Kransow, D. Nishita, M. Michel, D. V. Conti, J. Liu, W. Lee, C. K. Edlund, S. Hall, P. Y. Kwok, N. L. Benowitz, T. B. Baker, R. F. Tyndale, C. Lerman and G. E. Swan, "Nicotinic acetylcholine receptor variation and response to smoking cessation therapies", Pharmacogenetics and Genomics **23**, 2, 94–103 (2013).

Bickel, W. K., 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).

Boker, S. M., *APA Handbook of Research Methods in Psychology*, vol. 3 of *APA Handbooks in Psychology*, chap. Dynamical systems and differential equation models of change, pp. 323–333 (American Psychological Association, Washington, DC, 2012).

Boker, S. M. and J. Nesselroade, "A method for modeling the intrinsic dynamics of intraindividual variability: Recovering the parameters of simulated oscillators in multi-wave panel data", Multivariate Behavioral Research **37**, 1, 127–160 (2002).

Bollen, K. A., *Structural Equations with Latent Variables*, Wiley Series in Probability and Mathematical Statistics (John Wiley and Sons, New York, 1989).

Brdys, M. A., M. Grochowski, T. Gminski, K. Konarczak and M. Drewa, "Hierarchical predictive control of integrated wastewater treatment systems", Control Engineering Practice **16**, 751–767 (2008).

Brener, N. D., J. O. G. Billy and W. R. Grady, "Assessment of factors affecting the validity of self-reported health-risk behavior among adolescents: evidence from the scientific literature", Journal of Adolescent Health **33**, 6, 436–457 (2003).

Brier, M. E., A. E. Gaweda, A. Daily, G. R. Aronoff and A. A. Jacobs, "Randomized trial of model predictive control for improved anemia management", Clinical Journal of the American Society of Nephrology **5**, 5, 814–820 (2010).

Bullen, C., R. Whittaker, N. Walker and M. Wallace-Bell, "Pre-quitting nicotine replacement therapy: Findings from a pilot study", Tobacco Induced Diseases **3**, 35–40 (2006).

Camacho, E. F. and C. Bordons, *Model Predictive Control in the Process Industry*, Advances in Industrial Control (Springer-Verlag London Limited, London, England, 1995).

Camacho, E. F. and C. Bordons, *Model Predictive Control*, Advanced Textbooks in Control and Signal Processing (Springer, New York, NY, 1999).

Carpenter, M. J., B. F. Jardin, J. L. Burris, A. R. Mathew, R. A. Schnoll, N. A. Rigotti and K. M. Cummings, "Clinical strategies to enhance the efficacy of nicotine replacement therapy for smoking cessation: A review of the literature", Drugs **73**, 5, 407–426 (2013).

Carver, C. S. and M. F. Scheier, *On the Self-Regulation of Behavior* (Cambridge University Press, New York, New York, 1998).

Castro, Y., D. E. Kendzor, M. S. Businelle, C. A. Mazas, L. Cofta-Woerpel, P. M. Cinciripini and D. W. Wetter, "Structural and predictive equivalency of the Wisconsin Smoking Withdrawal Scale across three racial/ethnic groups", Nicotine and Tobacco Research **13**, 7, 548–555 (2011).

Centers for Disease Control and Prevention, "A practical guide to working with health-care systems on tobacco-use treatment", Tech. rep., U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health, Atlanta, GA (2006).

Centers for Disease Control and Prevention, "Vital signs: current cigarette smoking among adults aged $\geq$ 18 years − united states, 2009", MMWR **59**, 35, 1135–1140 (2010).

Centers for Disease Control and Prevention, "Current ciagette smoking among adults – united states, 2011", Morbidity and Mortality Weekly Report (MMWR) **61**, 44, 889–894 (2012).

Centers for Disease Control and Prevention, "Smoking & tobacco use: Quitting smoking", URL `www.cdc.gov` (2014).

Chandra, S., D. Scharf and S. Shiffman, "Within-day temporal patterns of smoking, withdrawal symptoms, and craving", Drug and Alcohol Dependence **117**, 118–125 (2011).

Chen, L., T. Baker, R. Grucza, J. Wang, E. Johnson, N. Breslau, D. Hatsukami, S. Smith, N. Saccone, S. Saccone, J. Rice, A. Goate and L. Bierut, "Dissection of the phenotypic and genotypic associations with nicotinic dependence", Nicotine and Tobacco Research **14**, 425–433 (2012).

Chisci, L., P. Falugi and G. Zappa, "Gain-scheduling mpc of nonlinear systems", International Journal of Robust and Nonlinear Control **13**, 295–308 (2003).

Collins, L. M., "Analysis of longitudinal data: The integration of theoretical model, temporal design, and statistical model", Annual Review of Psychology **57**, 1, 505–528 (2006).

Collins, L. M., "Unpacking the black box: engineering more potent interventions to improve public health. evan g. and helen g. pattishall outstanding research achievement award lecture presented at penn state university", URL `http://methodology.psu.edu/media/people/collins/pattishall.pdf` (2012).

Collins, L. M., J. W. Graham and B. P. Flaherty, "An alternative framework for defining mediation", Multivariate Behavioral Research **33**, 2, 295–312 (1998).

Collins, L. M., S. A. Murphy and K. L. Bierman, "A conceptual framework for adaptive preventive interventions", Prevention Science **5**, 3, 185–196 (2004).

Cui, Y., W. Wen, C. J. Moriarty and R. S. Levine, "Risk factors and their effects on the dynamic process of smoking relapse among veteran smokers", Behaviour Research and Therapy **44**, 7, 967–981 (2006).

Danaei, G., E. L. Ding, D. Mozaffarian, B. Taylor, J. Rehm, C. Murray and M. Ezzati, "The preventable causes of death in the united states: Comparative risk assessment of dietary, lifestyle, and metabolic risk factors", PLOS Medicine **6**, 4, e1000058 (2009).

DeLongis, A., S. Folkman and R. S. Lazarus, "The impact of daily stress on health and mood: Psychological and social resources as mediators", Journal of Personality and Social Psychology **54**, 3, 486–495 (1988).

Denai, M. A., M. Mahfouf and J. J. Ross, "A hybrid hierarchical decision support system for cardiac surgical intensive care patients. part i. physiological modelling and decision support system design", Artifical Intelligence in Medicine **45**, 1, 35–52 (2009).

Deshpande, S., *Optimal input signal design for data-centric identification and control with applications to behavioral health and medicine*, Ph.D. thesis, School of Electrical, Computer and Energy Engineering, Ira A. Fulton Schools of Engineering, Arizona State University, Tempe, AZ (2014).

Deshpande, S., N. N. Nandola, D. E. Rivera and J. Younger, "A control engineering approach for designing an optimized treatment plan for fibromyalgia", 2011 American Control Conference, San Francisco, CA pp. 4798–4803 (2011).

Deshpande, S., N. N. Nandola, D. E. Rivera and J. W. Younger, "Optimized treatment of fibromyalgia using system identification and hybrid model predictive control", Control Engineering Practice (in press).

Deshpande, S. and D. E. Rivera, "Optimal input signal design for data-centric estimation methods", Proceedings of the 2013 American Control Conference, Washington, DC pp. 3930–3935 (2013).

Deshpande, S. and D. E. Rivera, "Constrained optimal input signal design for data-centric estimation methods", IEEE Transactions on Automatic Control **59**, 11, 2990–2995 (2014).

Deshpande, S., D. E. Rivera and J. Younger, "Towards patient-friendly input signal design for optimized pain treatment interventions", Proceedings of the 16th IFAC Symposium on System Identification, Brussels, Belgium pp. 1311–1316 (2012).

Deshpande, S., D. E. Rivera, J. W. Younger and N. N. Nandola, "A control systems engineering approach for adaptive behavioral interventions: Illustration with a fibromyalgia intervention", Translational Behavioral Medicine **4**, 3, 275–289 (2014).

DiClemente, C. C., J. O. Prochaska, S. K. Fairhurst, W. F. Velicer, M. M. Velasquez and J. S. Rossi, "The process of smoking cessation: An analysis of precontemplation, contemplation, and preparation stages of change", Journal of Consulting and Clinical Psychology **59**, 2, 295–304 (1991).

Dong, Y., S. Deshpande, D. E. Rivera, D. S. Downs and J. S. Savage, "Hybrid Model Predictive Control for sequential decision policies in adaptive behavioral interventions", Proceedings of the 2014 American Control Conference pp. 4198–4203 (2014).

Dong, Y., D. E. Rivera, D. S. Downs, J. S. Savage, D. M. Thomas and L. M. Collins, "Hybrid model predictive control for optimizing gestational weight gain behavioral interventions", Proceedings of the American Control Conference (ACC) 2013 pp. 1970–1975 (2013).

Dong, Y., D. E. Rivera, D. M. Thomas, J. E. Navarro-Barrientos, D. S. Downs, J. S. Savage and L. M. Collins, "A dynamical systems model for improving gestational weight gain behavioral interventions", 2012 American Control Conference, Montreal, Canada (2012).

dos Santos, P. L., S. Deshpande, D. E. Rivera, T.-P. A. Perdicoulis, J. A. Ramos and J. Younger, "Identification of affine linear parameter varying models for adaptive interventions in fibromyalgia treatment", Proceedings of the 2013 American Control Conference, Washington, D.C. pp. 1976–1981 (2013).

Doyle, F. J., W. Bequette, R. Middleton, B. Ogunnaike, B. Paden, R. S. Parker and M. Vidyasagar, *The Impact of Control Technology: Overview, Success Stories, and Research Challenges*, www.ieeecss.org Control in Biological Systems, pp. 57–67 (IEEE Control Systems Society, 2011).

Doyle, F. J., L. M. Huyett, J. B. Lee, H. C. Zisser and E. Dassau, "Closed-loop artifical pancreas systems: Engineering the algorithms", Diabetes Care **37**, 5, 1191–1197 (2014).

Dube, S. R., A. McClave, C. James, R. Caraballo, R. Kaufmann and T. Pechacek, "Vital signs: Current cigarette smoking among adults $\geq$ 18 years—United States, 2009", MMWR: Morbidity and Mortality Weekly Report **59**, 35, 1135–1140 (2010).

Ehlert, U. and R. Straub, "Physiological and emotional response to psychological stressors in psychiatric and psychosomatic disorders", Annals of the New York Academy of Sciences **851**, 477–486 (1998).

Erhardt, L., "Cigarette smoking: An undertreated risk factor for cardiovascular disease", Atherosclerosis **205**, 1, 23–32 (2009).

Felder, R. M. and R. W. Rousseau, *Elementary Principles of Chemical Processes* (Wiley, Hoboken, NJ, 2004), 3rd update edn.

Fish, J. H. and J. R. Bartholomew, "Cigarette smoking and cardiovascular disease", Current Cardiovascular Risk Reports **1**, 5, 384–390 (2007).

Fortmann, S. P., T. Rogers, K. Vranizan, W. L. Haskell, D. S. Solomon and J. W. Farquhar, "Indirect measures of cigarette use: expired-air carbon monoxide versus plasma thiocyanate", Preventive Medicine **13**, 1, 127–135 (1984).

Franklin, G. F., J. D. Powell and M. L. Workman, *Digital Control of Dynamic Systems* (Ellis-Kagle Press, Half Moon Bay, CA, 1998), 3rd edn.

Franksen, O. I., "Mathematical programming in economics by physical analogies. Part I: The analogy between engineering and economics", Simulation pp. 297–341 (1969).

Fraser, D., K. Kobinsky, S. S. Smith, J. Kramer, W. E. Theobald and T. B. Baker, "Five population-based interventions for smoking cessation: a MOST trial", Translational Behavioral Medicine **in press** (2014).

Fraser Health, "Closed-loop control of anesthesia: Controlled delivery of remifentanil and propofol (icontrol-rp), clinicaltrials.gov identifier nct01771263", URL `http://clinicaltrials.gov/show/NCT01771263` (2014).

Garnier, H., M. Gilson, T. Bastogne and M. Mensler, *Identification of Continuous-Time Models From Sampled Data*, chap. The CONTSID Toolbox: A Software Support for Data-based Continuous-time Modelling, Advances in Industrial Control (Springer-Verlag London Limited, London, England, 2008a), 249-287 edn.

Garnier, H., M. Gilson and V. Laurain, "The CONTSID toolbox for Matlab: extensions and latest developments", Proceedings of the 15th IFAC Symposium on System Identification pp. 735–740 (2009).

Garnier, H., L. Wang and P. C. Young, *Identification of Continuous-Time Models From Sampled Data*, chap. Direct Identification of Continuous-time Models from Sampled Data: Issues, Basic Solutions and Relevance, pp. 1–29, Advances in Industrial Control (Springer-Verlag London Limited, London, England, 2008b).

Garnier, H. and P. C. Young, "The advantages of direclty identifying continuous-time transfer function models", International Journal of Control **87**, 7, 1319–1338 (2014).

Gaweda, A. E., A. A. Jacobs, G. R. Aronoff and M. E. Brier, "Model predictive control of erythropoietin administration in the anemia of ESRD", American Journal of Kidney Diseases **51**, 1, 71–79 (2008).

Ginexi, E. M., W. Riley, A. A. Atienza and P. L. Mabry, "The promise of intensive longitudinal data capture for behavioral health research", Nicotine and Tobacco Research **16**, S2, S73–S75 (2014).

Goodwin, G. C., M. M. Seron and J. A. Dedona, *Constrainted Control and Estimation: An Optimisation Approach*, Communications and Control Engineering (Springer-Verlag London Limited, London, England, 2005).

Gwaltney, C. J., J. Metrik, C. W. Kahler and S. Shiffman, "Self-efficacy and smoking cessation: A meta-analysis", Psychology of Addictive Behaviors **23**, 1, 56–66 (2009).

Gwaltney, C. J., S. Shiffman, M. H. Balabanis and J. A. Paty, "Dynamic self-efficacy and outcome expectancies: Prediction of smoking lapse and relapse", Journal of Abnormal Psychology **114**, 4, 661–675 (2005).

Haley, N. J., C. M. Axelrad and K. A. Tilton, "Validation of self-reported smoking behavior: biochemical analyses of cotinine and thiocyanate", American Journal of Public Health **73**, 10, 1204–1207 (1983).

Hall, K. D., S. B. Heymsfield, J. W. Kemnitz, S. Klein, D. A. Schoeller and J. R. Speakman, "Energy balance and its components: implications for body weight regulation", The American Journal of Clinical Nutrition **95**, 984–994 (2012).

Hamburg, M. A. and F. S. Collins, "The path to personalized medicine", New England Journal of Medicine **363**, 4, 301–304 (2010).

Hekler, E. B., M. P. Bunam, N. Poothakandiyil, D. E. Rivera, J. M. Dzierzewski, A. A. Morgan, C. S. McCrae, B. L. Roberts, M. Marsiske and P. R. Giacobbi, "Exploring behavioral markers of long-term physical activity maintenance: A case study of system identification modeling within a behavioral intervention", Health Education & Behavior **40**, 1S, 51S–62S (2013a).

Hekler, E. B., P. Klasnja, V. Traver and M. Hendriks, "Realizing effective behavioral management of health: The metamorphosis of behavioral science methods", IEEE Pulse **4**, 5, 29–34 (2013b).

Heron, K. E. and J. M. Smyth, "Ecological momentary interventions: Incorporating mobile technology into psychosocial and health behaviour treatments", British Journal of Health Psychology **15**, 1–39 (2010).

Holmstrom, K., A. O. Goran and M. M. Edvall, *User's Guide for TOMLAB / CPLEX v12.1*, TOMLAB OPTIMIZATION, Seattle, WA (2009).

Horst, W. D. and S. H. Preskorn, "Mechanisms of action and clinical characteristics of three atypical antidepressants: venlafaxine and nefazodone and bupropion", Journal of Affective Disorders **51**, 3, 237–254 (1998).

Hyland, M. E., "Control theory interpretation of psychological mechanisms of depression: comparison and integration of several theories", Psychological Bulletin **102**, 1, 109–121 (1987).

Intille, S. S., C. Kukla, R. Farzanfar and W. Bakr, "Just-in-time technology to encourage incremental, dietary behavior change", AMIA Annual Symposium Proceedings p. 874 (2003).

Ionides, E., C. Breto and A. King, "Inference for nonlinear dynamical systems", Proceedings of the National Academy of Sciences of the United States of America **103**, 49, 18438–18443 (2006).

Irvin, J. E. and T. H. Brandon, "The increasing recalcitrance of smokers in clinical trials", Nicotine and Tobacco Research **2**, 1, 79–84 (2000).

Irvin, J. E., P. S. Hendricks and T. H. Brandon, "The increasing recalcitrance of smokers in clinical trials ii: Pharmacotherapy trials", Nicotine and Tobacco Research **5**, 1, 27–35 (2003).

Kasari, C., A. Kaiser, K. Goods, J. Nietfeld, P. Mathy, R. Landa, S. Murphy and D. Almirall, "Communication interventions for minimally verbal children with autism: A Sequential Multiple Assignment Randomized trial", Journal of the American Academy of Child & Adolescent Psychiatry **53**, 6, 635–646 (2014).

Kendall, P. C., ed., *Child and Adolescent Therapy: Cognitive-Behavioral Procedures* (The Guilford Press, 2006), third edn.

Khuder, S. A., S. Milz, T. Jordan, J. Price, K. Silvestri and P. Butler, "The impact of a smoking ban on hospital admissions for coronary heart disease", Preventive Medicine **45**, 1, 3–8 (2007).

Killeen, P. R., "Markov model of smoking cessation", PNAS **108 Supplement 3**, 37, 15549–15556 (2011).

Kumar, S., W. Nilsen, M. Pavel and M. Srivastava, "Mobile health: Revolutionizing healthcare through trans-disciplinary research", Computer **46**, 1, 28–35 (2013).

Lee, J. H. and Z. H. Yu, "Tuning of Model Predictive Controllers for robust performance", Computers and Chemical Engineering **18**, 1, 15–37 (1994).

Lexicomp, "Bupropion: Drug information", Electronic database (2014).

Li, Y., J. L. Bienias and D. A. Bennett, "Confounding in the estimation of mediation effects", Computational Statistics and Data Analysis **51**, 6, 3173–3186 (2007).

Ljung, L., *System Identification: Theory for the User* (Prentice-Hall, Inc., Englewood Cliffs, New Jersey 07632, 1999), 2nd edn.

Ljung, L., "Experiments with identification of continuous time models", Tech. rep., Automatic Control at Linkopings Universitet, SE-581 83 Linkoping, Sweden (2009).

Ljung, L., *System Identification Toolbox: For Use with Matlab* (The MathWorks Inc., Natick, MA, 2011).

Ljung, L. and T. Glad, *Modeling of Dynamic Systems* (Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1994).

Ljung, L. and R. Singh, "Version 8 of the matlab System Identification Toolbox", Proceedings of the 16th IFAC Symposium on System Identification, Brussels, Belgium pp. 1826–1831 (2012).

Loh, W. Y., M. E. Piper, T. R. Schlam, M. C. Fiore, S. S. Smith, D. E. Jorenby, J. W. Cook, D. M. Bolt and T. B. Baker, "Should all smokers use combination smoking cessation pharmacotherapy? using novel analytic methods to detect differential treatment effects over 8 weeks of pharmacotherapy", Nicotine and Tobacco Research **14**, 2, 131–141 (2012).

Lopez, A. D., C. D. Mathers, M. Ezzati, D. T. Jamison and C. Murray, "Global and regional burden of dease and risk factors, 2001: systemic analysis of poulation health data", The Lancet **367**, 9524, 1747–1757 (2006).

Lopez-Meyer, P., S. Tiffany, Y. Patil and E. Sazonov, "Monitoring of cigarette smoking using wearable sensors and support vector machines", IEEE Transactions on Biomedical Enginering **60**, 7, 1867–1872 (2013).

MacKinnon, D., *Introduction to Statistical Mediation Analysis*, Multivariate Applications (Lawrence Erlbaum Associates, New York, New York, 2008).

Marlatt, G. and J. Gordon, "Determinants of relapse: Implications for the maintenance of behaviora change", Issue 78, part 7 of technical report, University of Washington Alcoholism & Drug Abuse Institute, Washington, USA (1980).

Martín, C. A., D. E. Rivera, W. T. Riley, E. B. Hekler, M. P. Bunam, M. A. Adams and A. C. King, "A dynamical systems model of Social Cognitive Theory", Proceedings of the 2014 American Control Conference, Portland, OR pp. 2407–2412 (2014).

Center for Disease Control and Prevention, "The great american smokeout", URL `http://www.cdc.gov/Features/GreatAmericanSmokeout/` (2011).

McCarthy, D. E., T. M. Piasecki, M. C. Fiore and T. B. Baker, "Life before and after quitting smoking: an electronic diary study", Journal of Abnormal Psychology **115**, 3, 454–466 (2006).

McCarthy, D. E., T. M. Piasecki, D. L. Lawrence, D. E. Jorenby, S. Shiffman and T. B. Baker, "Psychological mediators of bupropion sustained-release treatment for smoking cessation", Addiction **103**, 9, 1521–1533 (2008a).

McCarthy, D. E., 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 (2008b).

McCleary, R. and R. A. Hay, *Applied Time Series Analysis for the Social Sciences* (Sage Publications, Beverly Hills, CA, 1980).

McDowall, D., R. McCleary, E. E. Meidinger and R. A. Hay, *Interrupted time series analysis* (Sage Publications, Beverly Hills, CA, 1980).

McGuffie, K. and A. Henderson-Sellers, *A Climate Modelling Primer*, chap. Chapter 3: Energy Balance Models, pp. 81–116 (John Wiley and Sons, Inc., Hoboken, NJ, 2005), 3rd edn.

Meyer, G. K., T. B. Baker, S. S. Smith, M. Fiore, L. Redmond, T. W. Bosworth, P. Remington, D. Ahrens and A. Christiansen, "How smokers are quitting", Action paper 3, insights: Smoking in wisconsin, University of Wisconsin Center for Tobaccor Research & Intervention, Madison, WI (2003).

Molenaar, P., "Dynamic assessment and adaptive optimization of the psychotherapeutic process", Behavioral Assessment **9**, 389–416 (1987).

Molenaar, P. and C. G. Campbell, "The new person-specific paradigm in psychology", Current Directions in Psychology **18**, 2, 112–117 (2009).

Molenaar, P. C., "A manifesto on psychology as idiographic science: bringing the person back into scientific psychology, this time forever", Measurement Interdisciplinary Research and Perspectives **2**, 201–218 (2004).

Morari, M. and E. Zafiriou, *Robust Process Control* (Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1989).

Moskowitz, D. S. and S. N. Young, "Ecological momentary assessment: what it is and why it is a method of the future in clinical psychopharmacology", Journal of Psychiatry and Neuroscience **31**, 1, 13–20 (2006).

Nandola, N. N. 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).

Navarro-Barrientos, J. E., D. E. Rivera and L. M. Collins, "A dynamical model for describing behavioural interventions for weight loss and body composition change", Mathematical and Computer Modelling of Dynamical Systems **17**, 2, 183–203 (2011).

Noble, S. L., "Control-theoretic scheduling of psychotherapy and pharmacotherapy for the treatment of post-traumatic stress disorder", IET Control Theory & Applications **8**, 13, 1196–1206 (2014).

Novara, C., T.-P. A. Perdicoulis, J. A. Ramos and D. E. Rivera, *Linear Parameter-Varying System Identification: New Developments and Trends*, vol. 14 of *Advanced Series in Electrical and Computer Engineering*, chap. Introduction (World Scientific, Hackensack, NJ, 2011).

Ogunnaike, B. A. and W. H. Ray, *Process Dynamics, Modeling, and Control* (Oxford University Press, New York, 1994).

Patrick, D. L., A. Cheadle, D. C. Thompson, P. Diehr, T. Koepsell and S. Kinne, "The validity of self-reported smoking: a review and meta-analysis", American Journal of Public Health **84**, 7, 1086–1093 (1994).

Pawlowski, A., J. L. Guzman, J. E. Normey-Rico and M. Berenguel, "A practical approach for Generalized Predictive Control within an event-based framework", Computers and Chemical Engineering **41**, 52–66 (2012).

Personalized Medicine Coalition, "The case for personalized medicine, 4th edition", http://www.personalizedmedicinecoalition.org/, 1710 Rhode Island Ave., NW, Washington, DC (2014).

Piasecki, T. M., "Relapse to smoking", Clinical Psychology Review **26**, 196–215 (2006).

Pipe, A. L., S. Papdakis and R. D. Reid, "The role of smoking cessation in the prevention of coronary artery disease", Current Atherosclerosis Reports **12**, 2, 145–150 (2010).

Piper, M. E. and J. J. Curtin, "Tobacco withdrawal and negative affect: an anlysis of initial emotional response intensity and voluntary emotion regulation", Journal of Abnormal Psychology **115**, 1, 96–102 (2006).

Piper, M. E., E. B. Federman, D. E. McCarthy, D. M. Bolt, S. S. Smith, M. C. Fiore and T. B. Baker, "Using mediational models to explore the nature of tobacco motivation and tobacco treatment effects", Journal of Abnormal Psychology **117**, 94–105 (2008).

Piper, M. E., S. S. Smith, T. R. Schlam, M. C. Fiore, D. E. Jorenby, D. Fraser and T. B. Baker, "A randomized placebo-controlled clinical trial of 5 smoking cessation pharmacotherapies", Archives of General Psychiatry **66**, 11, 1253–1262 (2009).

Prochaska, J. O. and C. C. DiClemente, "Stages and processes of self-change of smoking: Toward an integrative model of change", Journal of Consulting and Clinical Psychology **51**, 3, 390–395 (1983).

Raphael, K., "Recall bais: A proposal for assessment and control", International Journal of Epidemiology **16**, 2, 167–170 (1987).

Rennard, S. I., N. A. Rigotti and D. M. Daughton, "Pharmacotherapy for smoking cessation in adults", Electronic database (2014).

Reynolds, S., "Combination therapy most effective for helping smokers quit", URL `www.drugabuse.gov/news-events/nida-notes/` (2011).

Riley, W. T., 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).

Rivera, D. E., "Optimized behavioral interventions: What does system identification and control systems engineering have to offer?", 16th IFAC Symposium on System Identification, Brussels, Belgium pp. 882–893 (2012).

Rivera, D. E., 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).

Ross, J. J., M. A. Denai and M. Mahfouf, "A hybrid hierarchical decision support system for cardiac surgical intensive care patients. part ii. clinical implementation and evaluation", Artifical Intelligence in Medicine **45**, 1, 53–62 (2009).

Rossiter, J. A., *Model-Based Predictive Control: A Practical Approach*, Control Series (CRC Press, Boca Raton, FL, 2003).

Rust, J. and S. Golmbok, *Modern psychometrics: The science of psychological assessment*, International Library of Psychology (Routledge, London, England, 1999), 2nd edn.

Savage, J. S., D. S. Downs, Y. Dong and D. E. Rivera, "Control systems engineering for optimizing a prenatal weight gain intervention to regulate infant birth weight", American Journal of Public Health (in press).

Scattolini, R. and P. Colaneri, "Hierarchical Model Predictive Control", Proceedings of the 46th IEEE Conference on Decision and Control pp. 4803–4808 (2007).

Schwager, M., "Data-driven identification of group dynamics for motion prediction and control", Journal of Field Robotics **25**, 6-7, 305–324 (2008).

Schwartz, J. D. and D. E. Rivera, "A process control approach to tactical inventory management in production-inventory systems", International Journal of Production Economics **125**, 1, 111–124 (2010).

Schwartz, J. D., W. Wang and D. E. Rivera, "Simulation-based optimization of process control policies for inventory management in supply chains", Automatica **42**, 8, 1311–1320 (2006).

Scokaert, P. and D. W. Clarke, *Advances in Model-Based Predictive Control*, chap. Stability and feasibility in constrained predictive control, pp. 217–229 (Oxford University Press, New York, NY, 1994).

Shadish, W., T. D. Cook and D. T. Campbell, *Experimental and quasi-experimental designs for generalized causal inference* (Houghton Mifflin, Boston, MA, 2002), 2nd edn.

Shiffman, S., "Relapse following smoking cessation: a situational analysis", Journal of Consulting and Clinical Psychology **50**, 1, 71–86 (1982).

Shiffman, S., "Dynamic influences on smoking relapse process", Journal of Personality **73**, 1715–1748 (2005).

Shiffman, S., D. M. Scharf, W. G. Shadel, C. J. Gwaltney, Q. Dang and S. M. Paton, "Analyzing milestones in smoking cessation: Illustration in a nicotine patch trial in adult smokers", Journal of Consulting and Clinical Psychology **74**, 2, 276–285 (2006).

Shiffman, S., A. Stone and M. Hufford, "Ecological momentary assessment", Annual Review of Pscychology **4**, 1-32 (2008).

Shiffman, S. and A. J. Waters, "Negative affect and smoking lapses: A prospective analysis", Journal of Consulting and Clinical Psychology **72**, 2, 192–201 (2004).

Solomon, R. L., *Psychopathology: Experimental Models*, chap. An Opponent-Process Theory of Acquired Motivation: The Affective Dynamics of Addiction, pp. 66–103 (W.H. Freeman, San Francisco, CA, 1977).

Solomon, R. L. and J. D. Corbit, "An opponent-process theory of motivation: I. temporal dynamics of affect", Psychological Review **81**, 119–145 (1974).

Soltesz, K., *On Automation in Anesthesia*, Ph.D. thesis, Department of Automatic Control, Lund University, Lund, Sweden (2013).

Steinberg, M. B., A. C. Schmelzer, P. N. Lin and G. Garcia, "Smoking as a chronic disease", Current Cardiovascular Risk Reports **4**, 413–420 (2010).

Tan, X., M. P. Shiyko, R. Li, Y. Li and L. Dierker, "A time-varying effect model for intensive longitudinal data", Psychological Methods **17**, 61–77 (2012).

The MathWorks, I., "c2d: Convert model from continuous to discrete time", URL `http://www.mathworks.com/help/control/ref/c2d.html` (2014a).

The MathWorks, I., "filter: 1-d digital filter", (2014b).

The MathWorks, I., "merge (iddata): Merge data sets into iddata object", URL `http://www.mathworks.com/help/ident/ref/mergeiddata.html` (2014c).

The MathWorks, I., "Model predictive control computation", URL `http://www.mathworks.com/help/mpc/index.html` (2014d).

The MathWorks, I., "pem: Prediction error estimate of linear or nonlinear model", URL `http://www.mathworks.com/help/ident/ref/pem.html` (2014e).

The MathWorks, I., "procestOptions: Options for procest", URL `http://www.mathworks.com/help/ident/ref/procestoptions.html` (2014f).

The MathWorks, I., "Simulink", URL `www.mathworks.com/products/simulink/` (2014g).

The MathWorks, I., "stepinfo: Rise time, settling time, and other step response characteristics", URL `http://www.mathworks.com/help/control/ref/stepinfo.html` (2014h).

Timms, K. P., 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 pp. 786–791 (2012).

Timms, K. P., D. E. Rivera, L. M. Collins and M. E. Piper, "Control systems engineering for understanding and optimizing smoking cessation interventions", Proceedings of the 2013 American Control Conference pp. 1967–1972 (2013).

Timms, K. P., D. E. Rivera, L. M. Collins and M. E. Piper, "Continuous-time system identification of a smoking cessation intervention", International Journal of Control **87**, 7, 1423–1437 (2014a).

Timms, K. P., D. E. Rivera, L. M. Collins and M. E. Piper, "Dynamic modeling and system identification of mediated behavior change with a smoking cessation case study", Submitted to Multivariate Behavioral Research (2014b).

Timms, K. P., D. E. Rivera, L. M. Collins and M. E. Piper, "A dynamical systems approach to understanding self-regulation in smoking cessation behavior change", Nicotine and Tobacco Research **16**, Suppl. 2, S159–S168 (2014c).

Timms, K. P., D. E. Rivera, M. E. Piper and L. M. Collins, "A Hybrid Model Predictive Control strategy for optimizing a smoking cessation intervention", Proceedings of the 2014 American Control Conference pp. 2389–2394 (2014d).

Tobacco Use and Dependence Guideline Panel, "A clinical practice guideline for treating tobacco use and dependence: 2008 update", Tech. rep., U.S. Department of Health and Human Services, Rockville, MD (2008).

Todd, M., "Daily processes in stress and smoking: Effects of negative events, nicotine dependence, and gender", Psychology of Addictive Behaviors **18**, 1, 31–39 (2004).

Trail, J. B., L. M. Collins, D. E. Rivera, R. Li, M. E. Piper and T. B. Baker, "Functional data analysis for dynamical system identification of behavioral processes", Psychological Methods **19**, 2, 175–187 (2014).

Urbina, S., *Essentials of Psychological Testing* (Wiley, Hoboken, NJ, 2011).

U.S. Department of Health and Human Services and Centers for Disease Control and Prevention, "The health consequences of smoking—50 years of progress: A report of the surgeon general", Tech. rep., U.S. Department of Health and Human Services, Public Health Services, Office of the Surgeon General, Rockville, MD (2014).

Van Henten, E. J. and J. Bontsema, "Time-scale decomposition of an optimal control problem in greenhouse climate management", Control Engineering Practice **17**, 66–96 (2009).

Velicer, W. F. and J. L. Fava, *Handbook of Psychology*, vol. 2, chap. Time Series Analysis for Psychological Research, pp. 581–606 (John Wiley and Sons, New York, NY, 2003).

Velicer, W. F., J. O. Prochaska, J. L. Fava, G. J. Norman and C. A. Redding, "Smoking cessation and stress management: Applications of the transtheoretical model of behavior change", Homeostasis in Health and Disease **35**, 5-6, 216–233 (1998).

Velicer, W. F., 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).

Walls, T. A. and D. E. Rivera, "Control engineering-based approaches to modeling substance abuse data", (17th Annual Meeting of the Society for Prevention Research, Washington, D.C., 2009).

Walls, T. A. and J. L. Schafer, *Models for Intensive Longitudinal Data* (Oxford University Press, Oxford, UK, 2006).

Wang, L. and H. Garnier, *System Identification, Environmental Modelling, and Control System Design* (Springer, New York, NY, 2012).

Warner, C. and M. Shoaib, "How does bupropion work as a smoking cessation aid?", Addiction Biology **10**, 219–231 (2005).

Watson, D. and L. A. Clark, *The PANAS-X: Manual for the Positive and Negative Affect Schedule - Expanded Form*, The University of Iowa, Iowa City, Iowa, 1999 update edn. (1999).

Watson, D., L. E. Clark and A. Tellegen, "Development and validation of brief measures of positive and negative affect: The panas scales", Journal of Personality and Social Psychology **54**, 6, 1063–1070 (1988).

Welsch, S. K., S. S. Smith, D. W. Wetter, D. E. Jorenby, M. C. Fiore and T. B. Baker, "Development and validation of the Wisconsin Smoking Withdrawal Scale", Experimental and Clinical Psychopharmacology **7**, 4, 354–361 (1999).

West, R., 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).

Widanage, W. D., K. R. Godfrey and A. H. Tan, "System identification in the presence of a saturation nonlinearity", Proceedings of the IFAC Workshop on Adaptation and Learning in Control and Signal Processing and IFAC Workshop on Periodic Control Systems (2004).

World Health Organization, "Preventing chronic diseases: A vital investment", Tech. rep., Department of Chronic Diseases and Health Promotion, World Health Organization, Geneva, Switzerland (2005).

World Health Organization, "The top 10 causes of death", URL `http://www.who.int/mediacentre/factsheets/fs310/en/` (2014).

Yasin, S. M., F. M. Moy, M. Retneswari, M. Isahak and D. Koh, "Timing and risk factors associated with relapse among smokers attempting to quit in Malaysia", International Journal of Tuberculosis and Lung Disease **16**, 7, 980–985 (2012).

Young, P. C., H. Garnier and M. Gilson, *Identification of Continuous-Time Models From Sampled Data*, chap. Refined instrumental variable identification of continuous-time hybrid Box-Jenkins models, Advances in Industrial Control (Springer-Verlag London Limited, London, England, 2008), 91-131 edn.

Zheng, A. and M. Morari, "Stability of model predictive control with mixed constraints", IEEE Transactions on Automatic Control **40**, 10, 1818–1823 (1995).

Zurakowski, R. and A. R. Teel, "A model predictive control based scheduling method for HIV therapy", Journal of Theoretical biology **238**, 368–382 (2006).

APPENDIX A

CUSTOM GUI FOR MODEL ESTIMATION & ANALYSIS

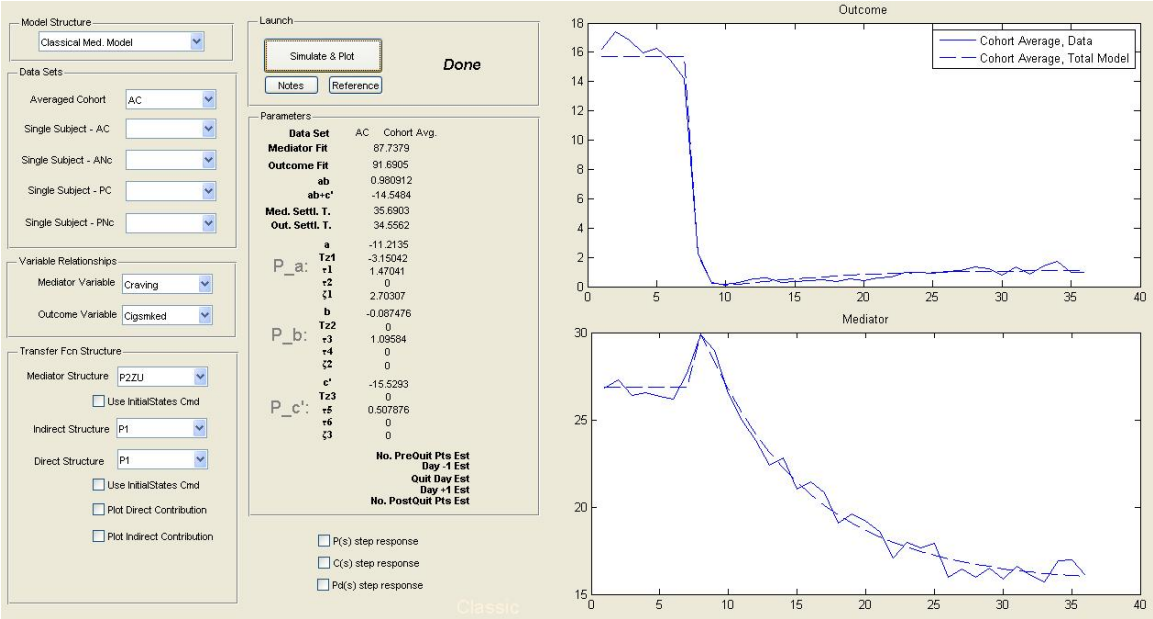

Figure A.1: Screenshot of the Custom Graphical User Interface Created in MATLAB for Flexible Analysis of Estimated Mediation and Self-Regulation Models Drawing from Group Average and Single Subject ILD from the McCarthy *et al.* (2008b) Study.