Available online at www.sciencedirect.com
ScienceDirect
Drug and Alcohol Dependence 88S (2007) S31–S40
DRUG and ALCOHOL DEPENDENCE
www.elsevier.com/locate/drugalcdep

Using engineering control principles to inform the design of adaptive interventions: A conceptual introduction

Daniel E. Rivera [a,*], Michael D. Pew [a], Linda M. Collins [b]

[a] *Control Systems Engineering Laboratory, Department of Chemical Engineering, Arizona State University Tempe, AZ 85287-6006, United States*
[b] *The Methodology Center and Department of Human Development and Family Studies, Penn State University, State College, PA 16801, United States*

Received 5 April 2006; received in revised form 12 October 2006; accepted 25 October 2006

Abstract

The goal of this paper is to describe the role that control engineering principles can play in developing and improving the efficacy of adaptive, time-varying interventions. It is demonstrated that adaptive interventions constitute a form of feedback control system in the context of behavioral health. Consequently, drawing from ideas in control engineering has the potential to significantly inform the analysis, design, and implementation of adaptive interventions, leading to improved adherence, better management of limited resources, a reduction of negative effects, and overall more effective interventions. This article illustrates how to express an adaptive intervention in control engineering terms, and how to use this framework in a computer simulation to investigate the anticipated impact of intervention design choices on efficacy. The potential benefits of operationalizing decision rules based on control engineering principles are particularly significant for adaptive interventions that involve multiple components or address co-morbidities, situations that pose significant challenges to conventional clinical practice.

Keywords: Adaptive interventions; Engineering process control; Substance abuse prevention

1. Introduction

Adaptive interventions represent a promising approach to prevention and treatment. They are especially useful for prevention programs with numerous components aimed at different aspects of risk, and for treatment of chronic, relapsing disorders such as alcoholism, cigarette smoking, and other types of substance abuse. Contingency management, individualized treatments, stepped care programs, and case management all represent frameworks that enable the implementation of adaptive interventions. Adaptive interventions individualize therapy by the use of *decision rules*, which express how the therapy level and type should vary according to *tailoring variables* such as response to treatment, adherence, and treatment burden (Murphy et al., 2007; MC-DATS, 2004). Adaptive interventions differ from conventional fixed interventions in significant ways. In fixed interventions, the same dosage is applied to all program participants without taking into account any of their individual characteristics. In an adaptive intervention, different dosages of prevention or treatment components are assigned to different individuals and/or to the same individual across time, with dosage varying in response to the needs of the individual. For example, the composition of a computer-delivered drug abuse prevention program might be varied somewhat depending on the ethnicity of the recipient. Adaptive interventions are time varying when the adaptation is repeated throughout the intervention. For example, a smoking cessation program may periodically assess each participant's progress along the stages of the Transtheoretical Model (Velicer and Prochaska, 1999), and accordingly adjust how key components of the intervention are presented. Adaptive interventions are strikingly similar to sensible clinical practice, but in order to be successful, they must be much more tightly managed than typical clinical procedures. Interest in adaptive techniques is significant not only in the treatment of substance abuse (Sobell and Sobell, 1999; Velicer and Prochaska, 1999; Brooner and Kidorf, 2002; Murphy and McKay, Winter 2003/Spring 2004) but also in the treatment of hypertension (Glasgow et al., 1989), depression (Rush et al., 2004), Alzheimer's disease (Schneider et al., 2001) and infectious diseases (Rosenberg et al., 2000).

Collins et al. (2004) articulated a general conceptual framework for the design of an adaptive intervention; however, their

* Corresponding author. Tel.: +1 480 965 9476; fax: +1 480 965 0037.
E-mail address: daniel.rivera@asu.edu (D.E. Rivera).

0376-8716/$ – see front matter © 2006 Elsevier Ireland Ltd. All rights reserved.
doi:10.1016/j.drugalcdep.2006.10.020

article did not address strategies for arriving at an intervention design that will produce the desired prevention and/or treatment outcomes. As adaptive strategies play an increasingly prominent role in addressing many important prevention and treatment problems, it is evident that much research is needed on adaptive intervention design and implementation. In this article we propose that ideas from control engineering can provide a basis for improving adaptive intervention design.

Control engineering is a broadly applicable field that is part of everyday life. It refers to the discipline that examines how to influence a *dynamical system* in order to regulate it and thereby achieve more desirable outcomes. A dynamical system is a multivariate time-varying process, usually nonlinear in nature, in which changes to input variables (some of which can be manipulated) lead to changes in output variables that affect outcomes of interest. Many items that form part of our modern existence rely on well-designed control systems to enable safe, profitable, and environmentally friendly operation; cruise control in automobiles, the home thermostat, and the insulin pump are all examples of control engineering at work. In the last few decades, significant improvements in computing technology, increasing access and availability to information, and novel methods for measurement and actuation have enabled the extensive application of control engineering concepts to physical systems.

Engineering control principles are applicable as well to areas in the behavioral sciences that involve dynamical systems, like time-varying adaptive interventions. Consider a time-varying adaptive intervention aimed at substance use treatment. Substance use, which is the outcome, varies over time, and is influenced by numerous time-varying variables (e.g., stress). The adaptive intervention is one of these time-varying influences on substance use, and importantly, unlike the other time-varying influences, it can be manipulated. The question is how to choose an adaptive intervention design so as to optimize the outcome. Control systems engineers routinely model complex dynamical systems and then apply control design algorithms and computer simulations based on these models to help them determine how to optimize outcomes. The same approach can be used to model and improve an adaptive intervention process. The control design and subsequent computer simulation tasks provide a means for experimenting with various decision rules, other design variables in the intervention, likely values of time-varying influences and various settings of characteristics that define the intervention participants, in order to investigate what will be the likely effects on key outcomes. The simulation results provide valuable information that can be used to choose decision rules and other aspects of the design so as to optimize the intervention. The resulting optimized intervention can then be evaluated in a randomized clinical trial.

The purpose of this article is to show some of the insights and benefits that potentially can be gained from a control engineering perspective on adaptive interventions, in the hope that intervention scientists will consider applying this perspective in their work. Section 2 describes some control engineering fundamentals and provides background useful in understanding the basic concepts in this field. Section 3 explores the link

Table 1
Fundamental control engineering and adaptive intervention terminology

Phrase	Definition
Block diagram	A graphical representation of the components that comprise a closed-loop control system.
Closed-loop	Refers to system behavior once a controller/decision policy is implemented.
Controller	A mathematical set of relationships that translate error (i.e., deviation from a goal or setpoint) into settings for a manipulated variable (which defines an intervention dosage). Also referred to as a decision policy or decision rule in the context of this report.
Control engineering	The science that considers how to manipulate system variables in order to transform dynamic behavior to desirable from undesirable.
Control error	The difference between the controlled variable and the goal or reference point; the ultimate goal of a control system is to take this variable to zero.
Control loop	Refers to a closed-loop system.
Controlled variables	System variables that we wish to keep at a reference value or setpoint.
Decision rules	Component of an adaptive intervention that decides on intervention dosages on the basis of values of a tailoring variable; a controller in engineering terminology.
Disturbance variable	A system input variable that influences the controlled variable response, but cannot be manipulated by the controller; disturbance changes occur external to the system (hence sometimes referred to as exogeneous variables).
Dynamical system	A multivariate time-varying process where changes in input variables (such as manipulated and disturbance variables) induces changes in outcomes of interest over time.
Feedback control	A control strategy in which a controlled variable is examined and compared to a reference value or setpoint. The controller issues actions (decisions on the values of a manipulated variable) on the basis of estimated values of the control error (the difference between the controlled variable measurement and the reference value).
Manipulated variable	A system input variable whose adjustment influences the response of the controlled variable; the magnitude of manipulated variable is determined by the controller.
Open-loop	Refers to dynamical system behavior without a controller or decision policy.
Offset	A sustained discrepancy between the controlled variable response and the setpoint in a closed-loop system, usually undesirable.
Process	The dynamical system under study, for which a closed-loop controller or decision rule will be applied.
Setpoint	Refers to a desired reference point or goal in the controlled variable which the control system is working to achieve.
Tailoring variable	A tailoring variable is a summary of available information that is used in an adaptive intervention context to determine the next course of treatment. In a feedback control context, the tailoring variable can act as a controlled variable, which needs to reach a desired goal at the conclusion of an intervention.

between adaptive, time-varying interventions and engineering control principles by examining a hypothetical intervention. A schematic representation of the hypothetical intervention in control engineering terms is presented, which gives rise to some important questions. The schematic representation enables the use of a simulation study as a tool to evaluate the intervention's decision rules under a variety of conditions, using a control engineering framework. Section 4 briefly describes a control engineering-based alternative to the decision rules, and discusses some systems and control technologies that are relevant to future research on this problem. Section 5 summarizes the major findings and observations of this article.

2. Introduction to control engineering principles

A vast number of introductory texts on control engineering are available; some examples include Seborg et al. (2004); Ogunnaike and Ray (1994), and Brosilow and Joseph (2002). Powers (1992) and Ramsay (2006) discuss control engineering concepts with a behavioral science focus in mind. Molenaar (1987) considers the use of feedback control to optimize the psychotherapeutic process. Table 1 presents a list of definitions that will be useful in understanding this topic, and which will be utilized throughout this paper.

Of the diverse forms of control engineering in the literature, *process control* holds significant promise for application in the design of adaptive behavioral interventions. Process control systems are widely used in the chemical industries to adjust flows to maintain level and product compositions at desired values (Ogunnaike and Ray, 1994). Effective adaptive interventions have goals similar to those of well-designed process control systems, in that both seek to (1) reduce negative effects, (2) increase intervention potency and (3) reduce waste (Collins et al., 2004). In this paper, the process control problem associated with maintaining liquid level in a tank (a.k.a., a fluid analogy) is proposed for modeling the dynamic behavior of adaptive interventions, and this class of models serves as the basis for decision policies evaluated using control engineering principles.

Feedback control represents one of the most useful and commonplace forms of control strategies applied in industrial practice. We can illustrate the concept of a feedback control system with a simple example from everyday life: taking a shower. Think of your ideal shower, with your preferred setting for temperature. Suppose a person in the shower controls temperature (see Fig. 1) by adjusting the hot water valve. In control theory language, the individual is the *controller*, temperature is the *controlled* variable, hot water flow is the *manipulated* variable and the particular temperature desired is the *setpoint* value. *Disturbance* (or exogenous) variables induce changes in the controlled variables that keep these from attaining the setpoint values. There are many possible sources of disturbances in this example, among them ambient temperature changes, fluctuations in the operation of the home water heater, and abrupt changes in the water flow sources to the shower (for instance, yard sprinklers going off or a nearby toilet being flushed). The purpose of the feedback control system, then, is to keep the controlled variable as close as possible to the setpoint value in spite

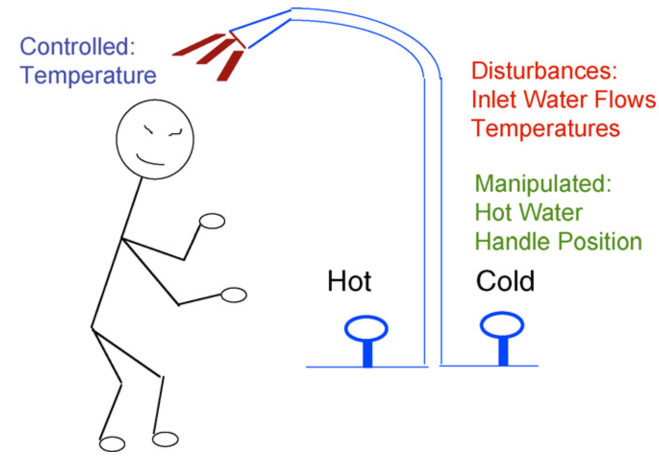

Fig. 1. Person in the shower: an everyday control problem.

of these disturbances, or in this case, to keep the temperature as close as possible to desired (ideal) settings. Manual feedback control of a shower as described here mimics conventional clinical practice, in that a clinician decides on dosages and treatment based on his/her judgement of the current state of the patient.

The action of the control system can be conducted continuously over time (in an analog controller) or at regular time intervals (in a sampled-data or digital controller). In either scenario, the controller makes necessary adjustments to the manipulated variables until the *control error* (the difference between the controlled variable values and setpoints) is minimized. The behavior of the control error in terms of temperature deviation for the shower problem, before and after a controller is enabled, is depicted in Fig. 2. Comparing the top and bottom panels of this figure helps illustrate a key concept in engineering control. Before the controller is enabled at $t = 2000$ there is considerable variability in temperature and no variability in hot water valve adjustment. However, once the controller has been enabled the situation is reversed; there is now little variability in temperature and considerable variability in hot water valve

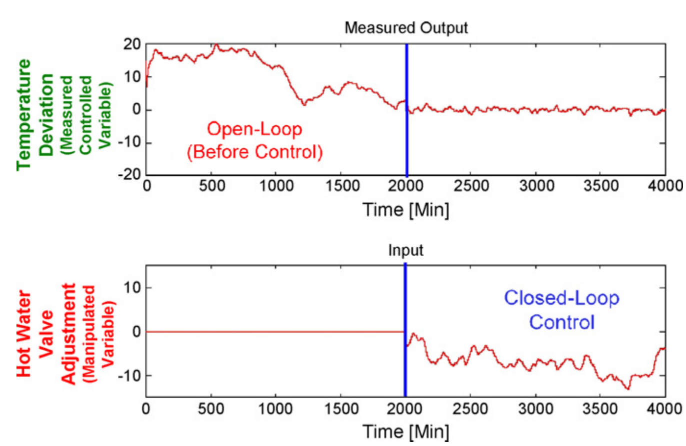

Fig. 2. Time series representing the effects of continuous feedback control action for the shower problem. From $t = 0$ to 2000, the system is in open-loop (manual) operation. Closed-loop (automatic) feedback control is engaged at $t = 2000$. The beneficial action of the control system is evidenced in the transfer of variability from temperature (the controlled variable) to the hot water valve position (the manipulated variable).

Table 2
Engineering control variables for the hypothetical *Fast Track* adaptive intervention

Adaptive intervention variable	Engineering control term	Notation
Intervention: dose of home visits	Manipulated variable	$I(t)$
Goal or threshold on parental function	Setpoint	$\mathrm{PF}^{\mathrm{Goal}}$
Extraneous sources depleting parental function	Disturbance input	$D(t)$
Tailoring variable: parental function	Controlled variable	$\mathrm{PF}(t)$
Review interval: time between adjustments to intervention dose	Controller sampling time	T

adjustment. The decisions made by the controller effectively transfer the variability from the "expensive" resource (in this case temperature, the controlled variable) to a "cheaper" one (the change in hot water flow, or manipulated variable). This transfer of variance is critical to the central role that engineering-based control systems play in industrial practice, and has important consequences for the use of these ideas in adaptive interventions for substance abuse prevention and treatment.

3. Adaptive interventions as engineering control systems

3.1. A hypothetical intervention

In this section we present a hypothetical adaptive intervention that will serve as the basis for a number of example problems described in this paper. The intervention is loosely based on the *Fast Track* program (CPPRG, 1992, 1999a, 1999b), which featured a number of adaptive components. The long-range purpose of the hypothetical intervention is to prevent the development of conduct disorders in children. The intervention component that is being delivered adaptively is family counseling. There are several possible levels of intensity, or doses, of family counseling. The idea is to vary the doses of family counseling depending on the need of the family, in order to avoid both providing an insufficient amount of counseling for very troubled families and wasting counseling resources on families that do not need it. The decision about which dose of counseling to offer each family is based on two factors. One is the family's level of functioning, assessed by a family functioning questionnaire completed by one of the parents. The score on the family functioning questionnaire is the tailoring variable, because it is used to determine the particular level of treatment provided to the individual family. The other factor is the judgment of a clinician familiar with the family's case. Based on the questionnaire and the clinician's assessment, family functioning is determined to fall in one of the following categories: very poor, poor, near threshold, or at/above threshold. The decision rule is as follows: families with very poor functioning are given weekly counseling; families with poor functioning are given biweekly counseling; families with near threshold functioning are given monthly counseling; and families at or above threshold are given no counseling. Family functioning is reassessed at a review interval of 3 months, at

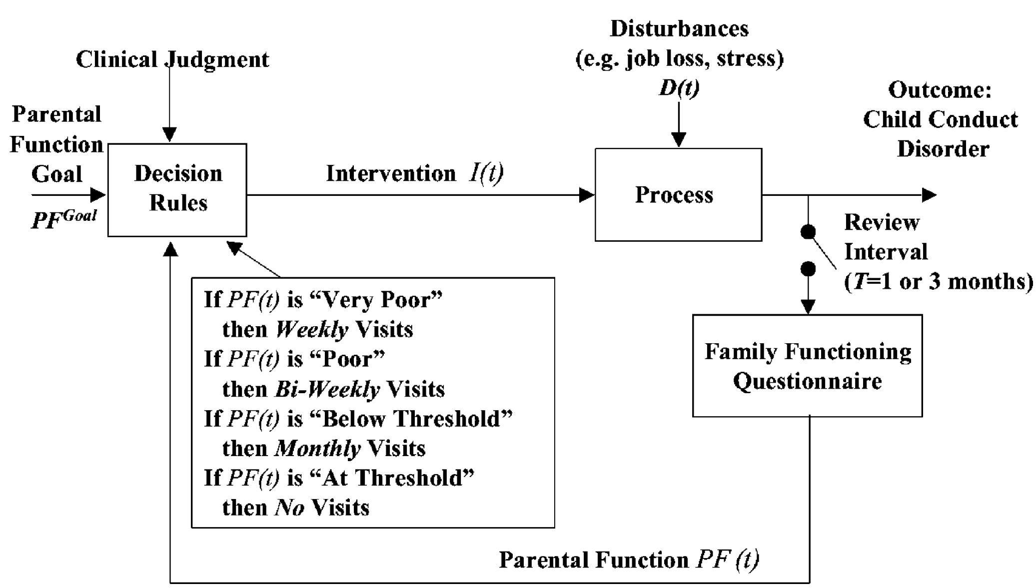

Fig. 3. Block diagram feedback control representation of the hypothetical time-varying adaptive intervention described in the text. In this representation, the measure of the tailoring variable, parental function, $\mathrm{PF}(t)$, acts as a controlled variable. A current value of the parental function (as assessed by a family functioning questionnaire) is received by the decision rules, which act as the feedback controller. The decision rules compare $\mathrm{PF}(t)$ to the goal $\mathrm{PF}^{\mathrm{Goal}}$ and determine a recommended dosage level for the intervention (in this case, the frequency of home visits). This recommendation is combined with clinical judgment to produce the final dosage assignment, denoted by $I(t)$. Once the intervention is implemented it acts upon the process that produces parental function. Also acting upon the process is the disturbance variable $D(t)$, which here represents the aggregate of time-varying characteristics or events that can deplete parental function (e.g., depression; loss of a job). At each successive review interval, the cycle is repeated; in other words, an updated value of the tailoring variable $\mathrm{PF}(t)$ is obtained and the feedback algorithm is performed once more. The final outcome of the intervention is the level of child conduct disorder.

which time the intervention dosage may change. This goes on for 3 years, with twelve opportunities for a dose of family counseling to be assigned. The final outcome of interest is a measure of conduct disorder in the target child, assessed 1 year after the end of the intervention period.

3.2. A block diagram representation of an adaptive intervention

This section describes how the adaptive intervention described in Section 3.1 can be cast as an engineering control system, specifically as a feedback control loop. This requires defining the various aspects of the adaptive intervention situation as engineering control variables, as is done in Table 2, and representing the adaptive intervention situation in a *block diagram* (Seborg et al., 2004; Ogunnaike and Ray, 1994), as is done in Fig. 3. In both Table 2 and Fig. 3 the time-varying nature of the variables is represented in the notation by (t). Because the adaptive intervention described above calls for assessment of family function at a review interval of 3 months, in this example time is discrete, and t represents review interval. For example, $I(4)$ represents the dose of intervention at review interval 4. (Continuous-time models are available for situations in which assessment is continuous or nearly so). Although the control-oriented representation of an adaptive intervention depicted in Fig. 3 is conceptual in nature, nonetheless it is extremely useful for articulating some fundamental questions on the design and implementation of adaptive, time-varying interventions. Among these questions are:

(1) *What role do disturbances* (*i.e.*, *individual time-varying characteristics*) *play in adaptive interventions*, *and to what extent can they be effectively managed by the actions of the decision policy*? Individual time-varying characteristics represent external (exogenous) conditions that influence how an individual (or a group of individuals) responds to an intervention. These "disturbances" form part of the control system and need to be effectively managed by the actions of a well-designed adaptive time-varying intervention. Specifically, the control system must suppress the deleterious effects of the detrimental disturbances, and take advantage of those that have a salutary effect and will help take the system to goal.
(2) *How frequently should the intervention dose be adjusted*? *In other words*, *what is the best review interval*? Previous work by Collins and Graham (2002) has demonstrated the importance of judicious selection of sampling time when drawing inference in longitudinal studies of substance use. In much the same way, the choice of review interval can be an important factor in determining the effectiveness of an adaptive intervention. In some situations shortening the review interval may not produce any appreciable gain in effectiveness; in others more frequent tailoring of the intervention may help the system approach goal more rapidly.
(3) *How can effective adaptive interventions be implemented in the presence of external clinical actions*? In many intervention settings it is the clinician who will ultimately make the final decision on dosage levels. An effective adaptive intervention should create a synergism between decision rules and clinical judgment; however, being able to accomplish this will depend in large part on the judicious choice of intervention design.
(4) *What decision rules will lead to the best outcomes*? Models that can describe prevention and treatment phenomena in a manner amenable to control engineering approaches can be used to arrive systematically at appropriate decision policies for problems related to these phenomena. A control engineering perspective dictates that the sophistication of the decision rules will be a function of the complexity of the model and the performance goals for the intervention.

Once an adaptive intervention has been recast as a feedback control system, as has been done in this section, it becomes possible to simulate the behavior of the system over time under any conditions that the investigator specifies. Such simulations are easily conducted using ordinary spreadsheet software, and can assist in providing answers to question like the ones above. This is illustrated in Section 3.3 below.

3.3. Simulated example of a control engineering approach to adaptive intervention design

3.3.1. Open-loop model definition. Fig. 3 describes the adaptive intervention as an engineering control loop. In order to conduct a computer simulation of an adaptive intervention, it is necessary to translate this pictorial representation into a mathematical model. The model is based on the fundamental scientific principle of conservation of matter, and can be conceptualized in terms of a *fluid analogy*, as depicted in Fig. 4. In Fig. 4, parental function corresponds to liquid material in a tank. The tank is depleted by time varying disturbances (outflow from the tank) and replenished by the intervention (inflow to the tank). Although particular disturbances could be beneficial in character (i.e., a new job, leaving a bad neighborhood, a positive faith-based experience, etc.), for purposes of this study we will treat the net sum of these disturbances as depletion. This analogy is similar to what is used to model production-inventory systems such as supply chains.

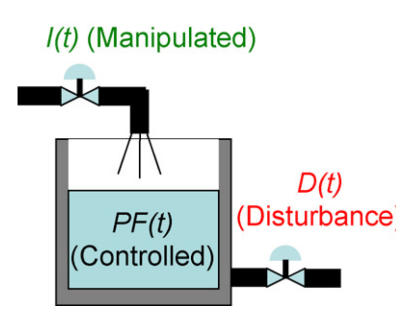

Fig. 4. Fluid analogy corresponding to the hypothetical adaptive intervention considered in the paper. Parental function PF(t) is treated as material (inventory) in a tank, which is depleted by disturbances $D(t)$, and replenished by intervention dosages $I(t)$, which is the manipulated variable.

In engineering systems, conservation and accounting of extensive properties such as mass, energy, momentum, and electrical charge serve as the basis for developing models that describe time-varying dynamical system behavior (Ogunnaike and Ray, 1994). The general accounting principle is represented by the equation:

$$\text{Accumulation} = \text{Inflow} - \text{Outflow} + \text{Generation} - \text{Consumption}, \quad (1)$$

In other words, inflow and generation are on the positive side, and outflow and consumption are on the negative side. Eq. (1) represents one possible approach to describe the "open-loop" dynamics of phenomena occurring in behavioral health problems.

Now that we have depicted the hypothetical adaptive intervention as a fluid analogy in Fig. 4, let us be more specific about expressing it in equation form. Consider a hypothetical family that is to participate in the example intervention, and also consider a time interval starting at time t and lasting T, so that the interval ends at time $t+T$. T is referred to as the review interval. In our hypothetical example, PF(t), which represents parental functioning at the start of the interval, adds to the level of parental functioning in the "tank." The intervention $I(t)$ also adds to the level in the "tank." Because families are likely to differ in how much they will benefit from an intervention, we include an intervention gain parameter, K_I, which represents how much this particular family benefits from intervention dose $I(t)$. Depleting the level of parental functioning in the "tank," as noted earlier and illustrated in Fig. 3, are numerous factors such as loss of a job, illness, and so on. In this example, for simplicity these depletion factors are summed to form a single disturbance parameter $D(t)$, which is a collective effect of multiple (n_d) individual time-varying characteristics.

Then the hypothetical intervention can be expressed in terms of the following difference equation, which models the relation between home visits and parental function:

$$\mathrm{PF}(t+T) = \mathrm{PF}(t) + K_I I(t) - D(t) \quad (2)$$

$$D(t) = \sum_{i=1}^{\mathrm{n_d}} D_i(t) \quad (3)$$

The Eq. (2) states that the parental function at the end of a review period (PF($t+T$)) equals the parental function at the start of the review period (PF(t)) plus the scaled effect of an intervention dose ($K_I\, I(t)$) less the depletion occuring during that time period ($D(t)$).

As noted earlier in this article, Eqs. (2) and (3) can be implemented in a computer program and used to simulate the parental function PF(t) response that would be expected over time under various forms and values of $D(t)$ and K_I. In conceptual terms, each computer simulation starts by specifying a hypothetical intervention participant, or subject. The simulation models what would be expected to happen if this hypothetical subject received the adaptive intervention over a particular period of time. The subject's values on any outcome variables of interest are tracked across the entire time period. In this way it is possible to examine whether changing aspects of the adaptive intervention, such as intervention dose or the frequency with which the dose is changed, are expected to make a difference in the outcome. As will be shown, it is instructive to repeat the simulation using hypothetical intervention participants with various characteristics, and also to vary the assumptions that underlie the simulation. For example, the simulations can be used to evaluate what kind of outcomes are likely to be associated with different values of $I(t)$, the intervention dosage recommended by the decision rules, and T, the review interval. Simulations can also be used to experiment with values of $D(t)$ and K_I to see what their effect on the outcome is likely to be.

It is important to note that the model per (2) and (3) hints at some of the significant modeling challenges associated with dynamically modeling prevention phenomena. The mechanisms by which interventions translate into outcomes that define the tailoring variable will most likely be uncertain and nonlinear in nature; these can also be highly auto correlated, involving multiple lagged values of interventions and tailoring variable measurements. A model parameter such as K_I will correspond in practice to a random variable, whereas individual disturbance effects contributing to the depletion rate $D(t)$ will have both deterministic and random components. Similar problems involving nonlinearity and uncertainty in the dynamic response are also seen in supply chain management problems, where fluid analogies serve as the basis for generating decision policies relying on control engineering principles (Braun et al., 2003; Schwartz et al., 2006).

3.3.2. Simulation as a heuristic tool for examining decision rules. In this subsection, we show simulations of the decision rules described in Section 3.1 (henceforth referred to as the rule-based control policy or rule-based controller) under various conditions of interest using a control engineering perspective.

Both parental function PF(t) and intervention dosage $I(t)$ are considered as normalized measurements with values ranging from 0 to 100%. These rules are mathematically summarized as follows:

- If parental function is "Very Poor" ($0 \le \mathrm{PF}(t) \le \mathrm{PF}^{\text{Very Poor}}$) then the intervention dosage should correspond to weekly home visits ($I(t) = I^{\text{weekly}}$),
- If parental function is "Poor" ($\mathrm{PF}^{\text{Very Poor}} < \mathrm{PF}(t) \le \mathrm{PF}^{\text{Poor}}$) then the intervention dosage should correspond to bi-weekly home visits ($I(t) = I^{\text{biweekly}}$),
- If parental function is "Below Threshold" ($\mathrm{PF}^{\text{Poor}} < \mathrm{PF}(t) < \mathrm{PF}^{\text{Goal}}$) then the intervention dosage should correspond to monthly home visits ($I(t) = I^{\text{monthly}}$),

If parental function meets or exceeds the goal ($\mathrm{PF}(t) \ge \mathrm{PF}^{\text{Goal}}$) then the intervention dosage should correspond to no home visits ($I(t) = 0$). The simulated response to the adaptive intervention will depend on the values selected for the threshold parameters $\mathrm{PF}^{\text{Very Poor}}$, $\mathrm{PF}^{\text{Poor}}$ and $\mathrm{PF}^{\text{Goal}}$ as well as the intervention dose reflected in the values of I^{weekly}, I^{biweekly}, and I^{monthly}. Settings for the parental function thresholds in the simulation are set to $\mathrm{PF}^{\text{Very Poor}} = 16.7\%$,

Table 3
Low and high settings for intervention gain (K_I) and depletion ($D(t)$) parameters for simulated cases considered in this paper

Settings	K_I (gain)	$D(t)$ (depletion)
Low	0.05	1.0
High	0.15	2.0

$\mathrm{PF}^{\mathrm{Poor}} = 33\%$, and $\mathrm{PF}^{\mathrm{Goal}} = 50\%$; the intervention potency is assumed to be linearly scaled and is defined according to $I^{\mathrm{weekly}} = 100\%$, $I^{\mathrm{biweekly}} = 66.7\%$, and $I^{\mathrm{monthly}} = 33\%$. Initially, the intervention subject is considered to possess 0% parental function with the intervention dosage determined at $t = 1$ month and reviewed every 3 months thereafter for a 36 month total program.

Many diverse and interesting simulation scenarios can be developed for this system. In this paper, we present a series of illuminating scenarios based on two hypothetical families under treatment that are distinguished by two parameters. One parameter, the gain K_I, represents how much the family is expected to benefit from the intervention. The other parameter, the depletion rate $D(t)$, corresponds to how much parental function is lost as a result of disturbances. Both the gain and the depletion rate are assumed constant per review instance in these simulations. The "low" and "high" settings for these parameters appear in Table 3.

Fig. 5 demonstrates the response obtained by applying an implementation of the decision rules described above to the family characterized by low intervention gain under no disturbances (i.e., zero depletion). This represents an ideal scenario, and the corresponding simulated response (Fig. 5) matches what would be expected from an efficacious adaptive intervention. Initially the rules dictate an intervention dosage of weekly home visits, with the frequency of visits decreasing as the parental function of the family improves. At 16 months parental function reaches the goal, and remains there for the duration of the assigned time period for the intervention. Having achieved the goal, the decision rules indicate that there is no need for additional home visits, and the intervention for that particular family is concluded.

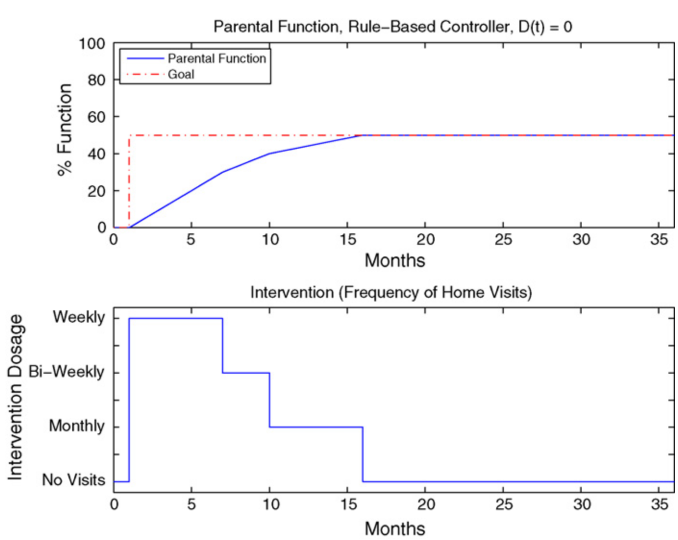

Fig. 5. Simulated closed-loop response of the decision rules of Section 3.3.2 for a family characterized by low intervention gain parameter K_I under no depletion. Review interval is quarterly.

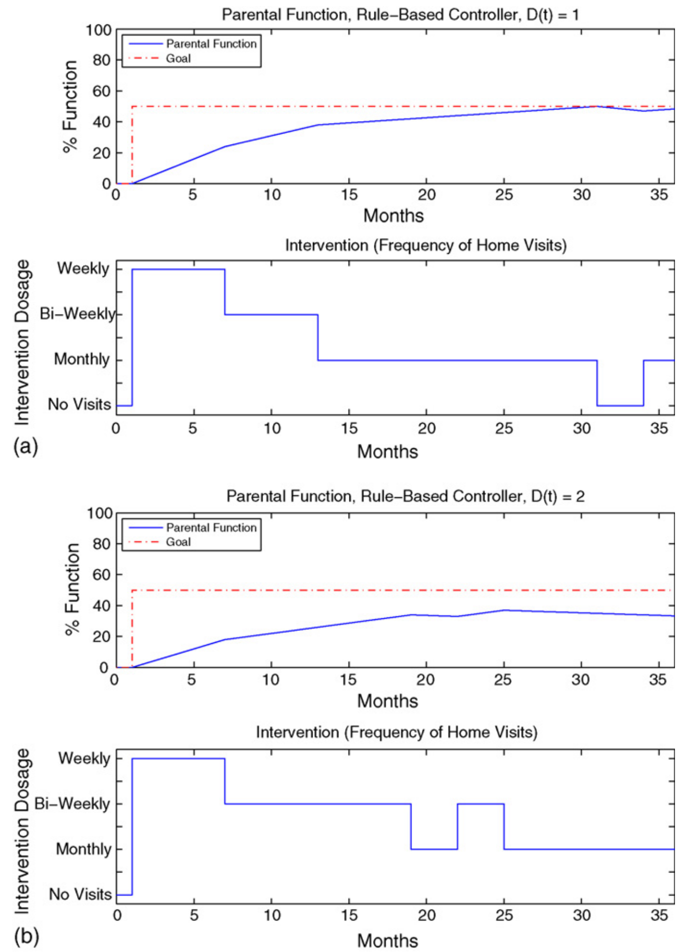

Fig. 6. Simulated closed-loop response of the decision rules of Section 3.3.2 for a family with low intervention gain under both low (a) and high (b) depletion rates $D(t)$. Review interval is quarterly.

Next we consider the presence of disturbances, in the form of depletion of parental function. The case where the rate of depletion is low is shown in Fig. 6a. Under these circumstances it takes longer for parental function to reach the goal (30 months) and the intervention dosages are on the average higher than before and never go to zero (since there is a constant loss of parental function per month). When the rate of depletion is high (Fig. 6b), the decision rules recommend higher dosages, but parental function fails to attain the goal throughout the 36 months of the intervention. This unattainment of the goal is referred to as "offset" in control engineering terminology. Despite a net increase in parental function in the family at the conclusion of the intervention, the resulting offset is an undesirable phenomenon that could be avoided by choosing an improved adaptive intervention design.

A simulation conducted for the family with the high setting for the intervention gain at this higher rate of depletion shows that the offset problem is not present (Fig. 7a). The higher gain value represents a family whose particular response to treatment results in a more effective translation of intervention dosage to positive outcomes, in comparison to the family with lower gain

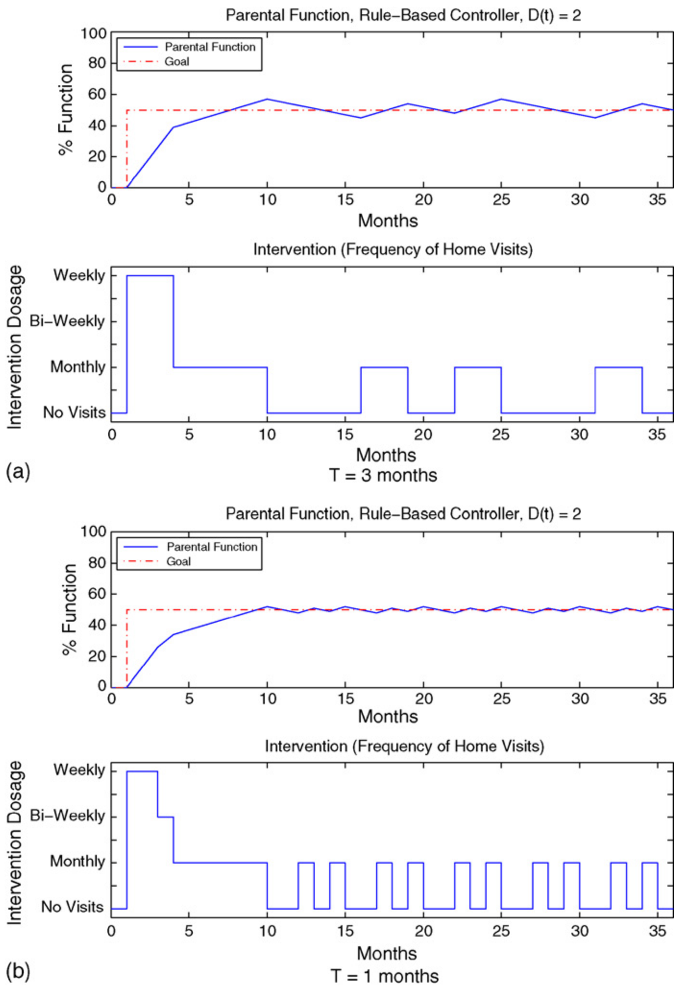

Fig. 7. Simulated closed-loop response of the decision rules of Section 3.3.2 for two values of the review interval T for a family with high intervention gain and high rate of depletion $D(t)$.

settings. For families displaying higher gains, the simulation results show that the use of these decision rules will lead to lower dosages at earlier stages in the intervention, while still achieving the desired outcomes. This simulation represents one scenario where adaptive delivery of intervention components results in reduced waste and improved allocation of resources.

Next we evaluate the effects of changing the review period T. The effect of reducing the review time interval T from 3 months to 1 month for the family with high gain and high depletion is presented in Fig. 7b. Comparing Fig. 7a and b shows that increasing the review frequency from quarterly to monthly does not really improve the overall effectiveness of the intervention. Although the intervention with monthly review intervals achieves increased tighter control of parental function around the goal than quarterly, the parental function goal is reached at roughly the same time frame ($\approx$9 months) at both review intervals. Furthermore, the shorter review period creates more frequent changes between dosage assignments (referred to as "chattering" in control engineering language) that may be viewed as undesirable to clinical personnel. Because applying a longer review period means lower costs, this simulation scenario is yet another example of how an engineering perspective can be used to optimize the delivery of an adaptive intervention.

Many additional interesting scenarios can be performed on this system, including ones involving increased levels of sophistication in both the open-loop dynamic model and the decision rules. For instance, Rivera et al. (2005) evaluate how a similar set of decision rules perform when the gain parameter K_I is not constant-valued, but changes exponentially with parental function, and the effects of introducing random noise in the parental function measurement. Rivera et al. (2005) furthermore describe how a decision rule based on a control-theoretic engineering design procedure can be obtained; some ideas on this topic are summarized in the ensuing section.

4. Decision rules based on engineering control principles

This section briefly describes how a model-based procedure for selecting optimized decision rules using control engineering principles can be applied to arrive at a decision policy for the simulated intervention that leads to improved outcomes. Some relevant control systems engineering technologies that can inform these optimized engineering-based interventions are also discussed.In general, the decision rules arising from control engineering methodologies will be defined by the character and sophistication of the open-loop model (i.e., Eqs. (2) and (3)), the desired performance characteristics expected for the intervention (e.g., how quickly should the parental function reach its goal, and what shape of response is desired for controlled and manipulated variables), and the types of disturbances that will be faced by the control system. Rivera et al. (2005) present a control law corresponding to the Proportional-Integral-Derivative (PID) family of controllers (Ogunnaike and Ray, 1994), which are among the most commonly applied control algorithms in industrial practice. In Rivera et al. (2005), a model-based approach referred to as Internal Model Control (Rivera et al., 1986; Morari and Zafiriou, 1989) is used to select parameters of a PID controller for a class of fluid analogy models that includes Eqs. (2) and (3). For the case of (2) and (3) the decision rule conforms to the equation

$$I(t) = I(t-T) + K_1 e(t) + K_2 e(t-T) \qquad (4)$$

where $e(t) = \mathrm{PF}^{\mathrm{Goal}} - \mathrm{PF}(t)$ is the deviation from goal or *control error* and K_1 and K_2 are controller parameters that are systematically determined from K_I and the character of $D(t)$. In the decision rule according to (4), the current dosage of the intervention is decided from the previous dosage recommendation, adjusted by scaled corrections from the current and previous control error.

The structure of (4) stands in sharp contrast to the "IF-THEN" decision rules summarized in Section 3.3.2, which assign dosages relying on only the most recent measurement of the tailoring variable. The dosage level computed in (4) is continuous in nature, and is assigned to the closest of one of four dosage levels (I^{weekly} I^{biweekly}, I^{monthly}, and 0) as described

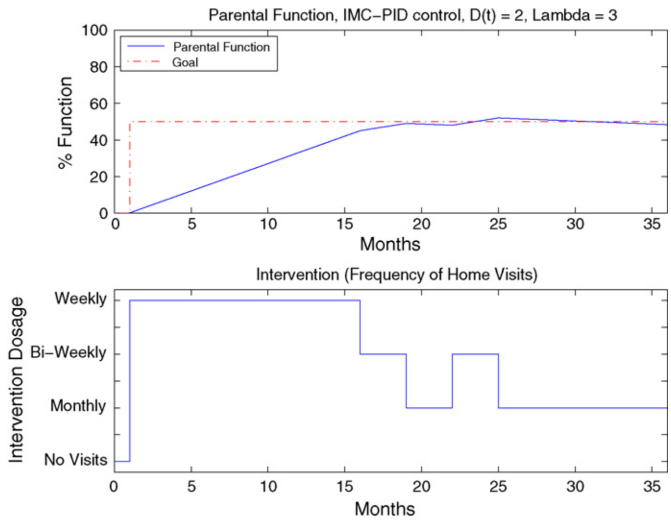

Fig. 8. Simulated closed-loop response of the engineering-based Proportional-Integral-Derivative (PID) control law tuned for the case of a family with low intervention gain under high depletion rate; the review interval is quarterly. The engineering-based decision rule eliminates the offset problem facing the decision rules of Section 3.3.2 under these same conditions by recommending the highest dosage over a longer period of time.

in Rivera et al. (2005). To see the impact that use of the more precise engineering-based decision rules can have on an adaptive intervention, recall the simulation results reported in Fig. 6b. According to this simulation, an intervention that used "IF-THEN" decision rules would fail to bring the level of parental functioning in this family, which has a low intervention gain and a high rate of depletion, up to the goal. Now consider Fig. 8, which shows the results of a simulation in which the same family participates in an intervention using engineering-based IMC–PID decision rules. The engineering-based rules eliminate the offset problem that plagued the "IF-THEN" rules, successfully bringing family functioning up to the goal level. Additional advantages of engineering-based decision rules are reviewed in Rivera et al. (2005).

The analysis of this paper motivates the study of two engineering technologies that may have substantial impact on adaptive interventions. One is obtaining open-loop dynamical models from experimental data via system identification, while the other is the synthesis of decision rules using Model Predictive Control. System identification refers to the field of study that is concerned with the modeling of dynamical systems from experimental data (Ljung and Glad, 1994; Ljung, 1999). In the system identification problem, data records of measured input and output variables are used to obtain a dynamical model that best approximates the phenomenon under study. System identification techniques have the potential to influence the development and dissemination of effective adaptive interventions in significant ways. In the context of the simulated example presented in this paper, system identification principles could be used to design an experimental trial that estimates the intervention gain parameter K_I and disturbance characteristics from patient data. It seems worthwhile to investigate how this field can influence current research efforts into modeling for adaptive interventions, for example, the use of sequential multiple assignment randomized trials (SMART) that have been proposed for generating datasets useful in the design of adaptive interventions (Collins et al., 2005; Murphy, 2005).

In this paper we presented an engineering-based PID decision policy that assigned dosages of one intervention component relying on values of one tailoring variable. For adaptive interventions involving multiple outcomes (such as those associated with co-morbidities) and multicomponent interventions, the concept of Model Predictive Control (MPC) (Prett and Garcia, 1988; Camacho and Bordons, 1999; Qin and Badgwell, 2003) seems promising. As a multivariable control design technique that uses optimization methods to make control decisions in real-time, MPC can serve as the basis for decision rules that involve multiple components and address multiple outcomes simultaneously while satisfying explicit problem constraints. Constraints can be imposed on the magnitudes and rate-of-change of intervention dosages, measured outcomes, and other system variables. Recent activity in using control engineering for adaptive interventions relies on MPC integrated with risk modeling approaches (Zafra-Cabeza et al., 2006) to take into account identified risks, their impact, and mitigating actions.

5. Summary and conclusions

This paper has established some conceptual linkages between the problem of adaptive interventions for prevention and treatment and engineering process control. Specifically, we have shown that adaptive, time-varying interventions are feedback control systems, with the outcome variable acting as the controlled variable, the intervention representing the manipulated variable, and decision rules serving the role of feedback control laws. The dynamics of the intervention in the "open-loop", that is, without decision rules, was represented using a fluid analogy corresponding to replenishing an inventory level in the face of depleting flows. A simulation study involving a hypothetical intervention inspired by the *Fast Track* program evaluated the performance of a control policy based on a series of IF-THEN decision rules under a variety of conditions. The simulations provided insights regarding the longitudinal trends (responses) observed in adaptive intervention for both tailoring variable and dosage assignments, potential problems such as offset (and how this is reflected in some families under treatment and not in others), the effect of time-varying individual characteristics (disturbances and gain parameters) and the benefits of an engineering-based decision rule based on PID control.

This simulated example, while providing some important insights, has not considered some significant challenges and issues: experimental trials that can estimate model parameters needed to design engineering-based decision rules, multivariable aspects (in terms of both potential multiple outcomes and multi-component interventions), and structural aspects (such as the need to address special populations with different rules). Exploring these issues is part of continuing research efforts in this problem.

Acknowledgement

Support for this work has been provided by the National Institute on Drug Abuse (NIDA) through a subcontract to Arizona State University (2778-ASU-DHHS-0075) from the Penn State University Methodology Center (prime award P50 DA010075).

Note: The Excel spreadsheet used to obtain the simulation results in this paper can be downloaded from the ASU Control Systems Engineering Laboratory website at: http://www.fulton.asu.edu/∼csel/Software.html.

References

Braun, M.W., Rivera, D.E., Carlyle, W.M., Kempf, K.G., 2003. Application of model predictive control to robust management of multi-echelon demand networks. Simul. Trans. Soc. Model. Simul. Int. 79 (3), 139–156.

Brooner, R.K., Kidorf, M., 2002. Using behavioral reinforcement to improve methadone treatment participation. Sci. Practice Perspect. 1, 38–48.

Brosilow, C., Joseph, B., 2002. Techniques of Model-Based Control. Prentice-Hall, Upper Saddle River, NJ.

Camacho, E.F., Bordons, C., 1999. Model Predictive Control. Springer-Verlag, London.

Collins, L.M., Graham, J., 2002. The effect of the timing and spacing of observations in longitudinal studies of tobacco and other drug use: temporal design considerations. Drug Alcohol Depend. 68, S85–S96.

Collins, L.M., Murphy, S.A., Bierman, K.L., 2004. A conceptual framework for adaptive preventive interventions. Prev. Sci. 5 (3), 185–196.

Collins, L.M., Murphy, S.A., Nair, V.N., Strecher, V.J., 2005. A strategy for optimizing and evaluating behavioral interventions. Ann. Behav. Med. 30, 65–73.

Conduct Problems Prevention Research Group, 1992. A developmental and clinical model for the prevention of conduct disorders: the fast track program. Dev. Psychopathol. 4, 509–527.

Conduct Problems Prevention Research Group, 1999a. Initial impact of the fast track prevention trial for conduct problems: I. The high-risk sample. J. Consult. Clin. Psychol. 67, 631–647.

Conduct Problems Prevention Research Group, 1999b. Initial impact of the fast track prevention trial for conduct problems: II. Classroom effects. J. Consult. Clin. Psychol. 67, 648–657.

Glasgow, M., Engel, B., D'Lugoff, B., 1989. A controlled study of a standardized behavioral stepped treatment for hypertension. Psychosom. Med. 51, 10–26.

Ljung, L., 1999. System Identification: Theory for the User, 2nd ed. Prentice-Hall, New Jersey.

Ljung, L., Glad, T., 1994. Modeling of Dynamic Systems. Prentice-Hall, New Jersey.

MC-DATS, 2004. Methodological challenges in developing adaptive treatment strategies. http://neuromancer.eecs.umich.edu/dtr/twiki/bin/view, accessed on September 18, 2006.

Molenaar, P., 1987. Dynamic assessment and adaptive optimization of the psychotherapeutic process. Behav. Assess. 9, 389–416.

Morari, M., Zafiriou, E., 1989. Robust Process Control. Prentice Hall, Englewood Cliffs, NJ.

Murphy, S.A., 2005. An experimental design for the development of adaptive treatment strategies. Stat. Med. 24, 1455–1481.

Murphy, S.A., McKay, J.R. Winter 2003/Spring 2004. Adaptive treatment strategies: an emerging approach for improving treatment effectiveness. Clinical Science (Newsletter of the American Psychological Association Division 12, section III: The Society for the Science of Clinical Psychology).

Murphy, S.A., Oslin, D.W., Rush, A.J., Zhu, J., 2007. Methodological challenges in constructing effective treatment sequences for chronic psychiatric disorders. Neuropsychopharmacology 32, 257–262.

Ogunnaike, B.A., Ray, W.H., 1994. Process Dynamics, Modeling, and Control. Oxford University Press, New York.

Powers, W.T., 1992. A cognitive control system. In: Levine, R.L., Fitzgerald, H. (Eds.), Analysis of Dynamic Psychological Systems, vol. 2: Methods and Applications. Plenum Press, New York, pp. 327–340.

Prett, D., Garcia, C., 1988. Fundamental Process Control. Butterworth, Stoneham, MA.

Qin, S.J., Badgwell, T.A., 2003. A survey of industrial model predictive control technology. Control Eng. Practice 11, 733–764.

Ramsay, J.O., 2006. The control of behavioral input-output systems. In: Walls, T., Schafer, J. (Eds.), Models for Intensive Longitudinal Data. Oxford University Press, New York, pp. 176–194.

Rivera, D.E., Morari, M., Skogestad, S., 1986. Internal model control. 4. PID controller design. Ind. Eng. Chem. Process Des. Dev. 25, 252–265.

Rivera, D.E., Pew, M.D., Collins, L.M., Murphy, S.A., 2005. Engineering control approaches for the design and analysis of adaptive, time-varying interventions. The Methodology Center Technical Report 05-73, Pennsylvania State University.

Rosenberg, E.S., Altfield, M., Poon, S.H., Phillips, M.N., Wilkes, B.M., Eldridge, R.L., Robbins, G.K., D'Aquila, R.T., Goulder, P., Walker, B.D., 2000. Immune control of HIV-1 after early treatment of acute infection. Nature 407, 523–526.

Rush, A.J., Fava, M., Wisniewski, S.R., Lavori, P.W., Trivedi, M.H., Sack-eim, H.A., Thase, M.E., Nierenberg, A.A., Quitkin, F.M., Kashner, T.M., Kupfer, D.J., Rosenbaum, J.F., Alpert, J., Stewart, J., McGrath, P.J., Biggs, M.M., Shores-Wilson, K., Lebowitz, B.D., Ritz, L., Niederehe, G., 2004. Sequenced treatment alternatives to relieve depression (STAR*D): rationale and design. Control. Clin. Trials 25 (1), 119–142.

Schneider, L., Tariot, P., Lyketsos, C., KS, K.D., Davis, K., Davis, S., Hsiao, J., Jeste, D., Katz, I., Olin, J., Pollock, B., Rabins, P., Small, G., Lebowitz, B., Lieberman, J., 2001. National Institute of Mental Health Clinical Antipsychotic Trials in Intervention Effectiveness (CATIE). Am. J. Geriatr. Psychiatry 9, 346–360.

Schwartz, J.D., Wang, W., Rivera, D.E., 2006. Simulation-based optimization of model predictive control policies for inventory management in supply chains. Automatica 42 (8), 1311–1320.

Seborg, D.E., Edgar, T.F., Mellichamp, D.A., 2004. Process Dynamics and Control, 2nd ed. Wiley, New York, NY.

Sobell, M.B., Sobell, L.C., 1999. Stepped care for alcohol problems: an efficient method for planning and delivering critical services. In: Tucker, J., Donovan, D., Marlatt, G. (Eds.), Changing Addictive Behavior: Bridging Clinical and Public Health Strategies. Guilford Press, New York, pp. 331–343.

Velicer, W.F., Prochaska, J.O., 1999. An expert system intervention for smoking cessation. Patient Educ. Couns. 36, 119–129.

Zafra-Cabeza, A., Rivera, D. E., Collins, L. M., Ridao, M. A., Camacho, E. F. 2006 A risk-based model predictive control approach to adaptive interventions in behavioral health. In: IEEE Conference Decision Control, pp. 673–678.