ON THE CONTROL OF PSYCHOLOGICAL NETWORKS

TEAGUE R. HENRY

DEPARTMENT OF PSYCHIATRY, UNIVERSITY OF PITTSBURGH

DONALD J. ROBINAUGH

DEPARTMENT OF PSYCHIATRY, HARVARD MEDICAL SCHOOL & MASSACHUSETTS
GENERAL HOSPITAL

EIKO I. FRIED

DEPARTMENT OF CLINICAL PSYCHOLOGY, LEIDEN UNIVERSITY

This project was supported by the American Foundation for Suicide Prevention, the Charles A. King Trust Postdoctoral Research Fellowship Program, Bank of America, N.A., Co-Trustees, and a National Institute of Mental Health Career Development Award (1K23MH113805-01A1) awarded to D. Robinaugh.

We would like to acknowledge the two anonymous reviewers and AE, whose comments and suggestions drastically improved this work, and Dr. Jennifer MacCormack for their helpful feedback on earlier drafts.

Abstract

The combination of network theory and network psychometric methods have opened up a variety of new ways to conceptualize and study psychological disorders. The idea of psychological disorders as dynamic systems has sparked interest in developing interventions based on results of network analytic tools. However, estimating a network model is not sufficient for determining which symptoms might be most effective to intervene upon, nor is it sufficient for determining the potential efficacy of any given intervention. In this paper, we attempt to remedy this gap by introducing fundamental concepts in control theory to both methodologists and applied psychologists. We introduce two controllability measures to the psychometric literature, average and modal controllability, to facilitate selecting the best set of intervention targets. Following this introduction, we (a) show how intervention scientists can probe the effects of both theoretical and empirical interventions on networks derived from real data; (b) demonstrate how simulations can account for intervention cost and the desire to reduce specific symptoms; (c) introduce the intervention efficacy ratio (IER), a metric for evaluating intervention efficacy; and (d) showcase how between subject heterogeneity in intervention response can be evaluated via a statistical testing procedure for determining if the dynamical systems of different people have the same optimal intervention targets. Every step is illustrated using rich clinical EMA data from a sample of subjects undergoing treatment for complicated grief, with a focus on the outcome *suicidal ideation*. All methods are implemented in an open-source R package `netcontrol`, and complete code for replicating the analyses in this manuscript is available at <https://osf.io/f268v/>

Key words: control theory; network psychometrics; complicated grief; personalized medicine; intervention science; psychopathology

Common cause models of psychopathology posit that correlations between symptoms arise because those symptoms stem from a shared underlying disease entity (e.g., coughing, fever, and Koplik spots co-occur due to the measles virus). Although this common cause model is clearly appropriate for some diseases, such as those caused by pathogens, the utility of this perspective for psychopathology has been increasingly questioned (Borsboom, 2008; Kendler, Zachar, & Craver, 2011). In contrast, the network theory of psychopathology posits that psychological disorders are sets of causally interacting symptoms (Borsboom, 2017; Cramer, Waldorp, van der Maas, & Borsboom, 2010). Psychopathology is an emergent property constructed from dynamic interactions among symptoms and other related problems (Kendler et al., 2011; Radden, 2018). This theoretical stance is grounded in insights drawn from philosophy of psychiatry, network science, systems science, and historical calls for greater complexity in clinical psychology and, in recent years, it has generated a rapidly growing body of empirical research (Robinaugh, Hoekstra, Toner, & Borsboom, 2019).

Researchers working within the network approach to psychopathology have relied heavily on *network psychometrics*: a class of statistical models derived from diverse fields (e.g., statistical mechanics) that seeks to model relations between symptoms (Epskamp, 2017). Researchers using these methods often focus their analyses on identifying "central" symptoms strongly interconnected within the symptom network. The hope is that such analyses may identify symptoms that play an important causal role in the maintenance of the disorder and, thus, may be an important target for treatment (Contreras, Nieto, Valiente, Espinosa, & Vazquez, 2019; Fried et al., 2017). Although early network studies often assessed cross-sectional (i.e., between-subject) data, there has been a growing focus on examining temporal (i.e., within-subject) data, as this level of analysis is better aligned with the within-person hypotheses about disease etiology and treatment that are the focus of network theory (Jordan, Winer, & Salem, 2020). It is posited that targeted treatment of "central" symptoms in these temporal networks (i.e., those exhibiting strong temporal associations with other symptoms) will produce considerable downstream improvement in other symptoms, a position that is consistent with how many clinicians think about mental disorder treatment (Kim & Ahn, 2002). Clinical trials testing this idea are already underway (Levine & Leucht, 2016).

Unfortunately, predicting how complex, dynamical systems will respond to intervention is exceedingly difficult and there is no guarantee that treatments targeting central symptoms will actually offer the hypothesized benefits outlined above. Recent work examining the relation between the centrality of a symptom and its causal influence found that the majority of graph theoretic centrality measures are not associated with the causal influence of that node (Dablander & Hinne, 2019). In this paper, we will propose that *control theory*, a well-established mathematical discipline, provides a better, data-driven way to identify optimal treatment targets while also equipping us to evaluate the theoretical efficacy of proposed interventions using

estimated network structures.¹ We will begin by introducing readers to network psychometrics and to the core concepts in control theory. After these primers, we introduce optimal control methods and demonstrate in two simulated examples how these control theory methods can be applied to psychological data to solve the inferential gap between identifying central nodes and determining intervention efficacy (Fried, 2020). We then further illustrate the clinical utility of control theory by using it in a series of simulation studies to (a) examine single target interventions, (b) evaluate multiple interventions applied simultaneously, (c) simulate real interventions, and (d) evaluate treatment response heterogeneity. We conclude by discussing the limitations of this approach and the steps needed in future research to fully realize the potential of control theory within psychological science.

1. Introduction to Control Theory

The first step to any application of control theory is that of *system identification* (Ljung, 1999). In traditional control theory settings such as manufacturing or engineering, the systems in question are often *a priori* known, such as a production line. In the case of psychopathology, the dynamics governing the relations between symptoms overtime are unknown, and must be determined. To do this, we construct a psychological network. A psychological network consists of *nodes*, representing the variables under study, and *edges*, representing some kind of relation between the variables, be it causal or statistical. Edges may represent the within-subject relations among variables within a single individual over time or the between-subject relations among variables in a group of individuals. Consider the network presented in Fig. 1 estimated using multilevel integrated vector autoregressive models (ML-VARI Bringmann et al., 2013; Epskamp, Waldorp, Möttus, & Borsboom, 2018; Shumway & Stoffer, 2017) using ecological momentary assessment (EMA) data of symptoms from 8 subjects during treatment for complicated grief. Complicated grief is a bereavement-specific syndrome characterized by persistent, intense, and impairing grief-related thoughts, emotions, and behaviors. Nodes in the network represent 9 core symptoms of complicated grief as well as suicidal ideation, or thoughts about suicide. Edges represent the relation between the change in a symptom from time $t - 1$ to t and the change in another symptom from time t to time $t + 1$, controlling for all other lagged relations. Because relations are temporal and differ in magnitude, the resulting edges are both *directed* and *weighted*. That is, they provide information about both the direction of the temporal association (e.g., change in thoughts about the future predicts change in emotional pain, but not vice versa) and the strength of that association (e.g., the association from suicidal ideation to thoughts about the loved one is stronger than the reverse).

¹We intend this manuscript to serve as an introduction to control theory both for methodologists working within clinical sciences as well as clinical psychologists wishing to better model disorders and possible treatment effects. Accordingly, although many technical details are included in the main text, some have been omitted to provide a more accessible overview. Further technical details about the models and code to replicate all results presented in this paper are included in the Supplementary Materials (SMs) and at <https://osf.io/f268v/>. All controllability centrality measures and optimal control inputs are calculated using `netcontrol` (Henry, 2020), an open-source, publicly available R package implementing many methods from control theory for use in psychological and neurological data analysis.

FIGURE 1.

Symptom Dynamics of Complicated Grief. Blue edges represent positive relations, red represents negative relations. Autoregressive effects were rescaled to improve edge contrast, and edges with absolute values below .05 were omitted from the visualization. Note that this network represents the lagged relation between **the changes in symptom values** (i.e., the use of an integrated VAR) rather than the lagged relations between the symptom values themselves. For details on what the integrated VAR model is and why it was chosen here, see SMs

Importantly, a network model-informed intervention requires more than just estimating network structure. As discussed above, the behavior of a given dynamical system, such as the ones proposed by network psychometric models for psychological disorders, can be difficult to predict on the basis of the network structure alone. An intervention does not have just an instantaneous effect, rather that instantaneous treatment effect will propagate over time throughout the dynamical system, which in turn can change the overall impact of the intervention. To understand the impact of interventions on a dynamical system, an approach that goes beyond structure to examine how the complex dynamics of the system will be affected over time is required are required. Control theory provides one such toolkit.

Control theory is a mathematical discipline concerned with the "control" of dynamic systems. It was first developed to improve manufacturing outcomes. For example, helping to determine the series of inputs to a manufacturing plant that optimize production while minimizing cost. Since its development, control theory has emerged as an important area of research in diverse areas of science. In this section, we will introduce several fundamental concepts of control theory with a focus on its application in psychopathology. To make this introduction as accessible as possible, we will primarily use terminology from psychology and network science and will include equivalent terms from control theory parenthetically and in bold italics (e.g., psychological term (***control theory term***)). To further assist translation, a glossary of control theory terms and their equivalent concepts in psychology appears in Table 1.

TABLE 1.
Glossary of Control Theory Terms, Psychological Equivalent, and Definition

Control Theory	Psychological Equivalent	Definition
Physical plant	Psychological system	The dynamical system that is being controlled
Control inputs	Interventions/treatments	Intentional inputs to the system
Disturbance	Disturbance	Random inputs to the system (not measurement error)
Error	Measurement Error	Random error due to measurement (doesn't impact the system)
State	Symptom value	The (possibly unobserved) values of the variables in the system
Observation	Measurement / Item	The observed measurement of a variable in the system
Energy	Cost of intervention	The overall cost of applying control inputs

Notably, we are not the first to propose applying control theory to psychological research. Hyland (1987) applied principles of control theory to theoretical models of depression, while Carver and Scheier (1998) and Johnson, Chang, and Lord (2006) examined how control theory principles map onto self-regulation, cognition, and behavior. Other researchers have applied control theory principles to adaptive intervention designs. For example, Molenaar (1987), Sinclair and Molenaar (2008) and Molenaar (2010) demonstrated how optimal control methods can be used to individualize psychotherapeutical trajectories. Rivera, Pew, and Collins (2007) likewise described how to apply control theory to model and assess adaptive interventions on single target variables. In perhaps the most comprehensive application, Rivera, Hekler, Savage, and Downs (2018) used a sophisticated form of control called *model predictive control* (MPC) in two separate interventions with sedentary, overweight adults and pregnant women (for further discussion of the utility of MPC, see our Future Directions section). Finally, there is a related large literature on *adaptive interventions* (Collins, Murphy, & Bierman, 2004; Collins et al., 2004; Nahum-Shani et al., 2018), which is considerably broader than control theory in intervention design, but addresses key considerations for both study and intervention design.

Prior applications of control theory have primarily focused on either theoretical implications or *single* target behavioral interventions (e.g., walking). With the development of network psychometrics and the expansion of intensive longitudinal data collection, there is a growing opportunity for more complex applications of control theory; applications which promise to not only tailor existing interventions to individual needs, but also to improve our understanding of psychological disorders by identifying both core controllable symptoms and heterogeneity in these symptoms and their response to treatment. Accordingly, our introduction to control theory will focus on the potential for application in these growing areas of research.

1.1. Defining The Control System

There are many possible system types, including continuous time, non-linear, time-varying, and stochastic systems. Here, it is important to note the distinction between a system in the control theory sense, and a model in the statistical sense. A model refers to a specific mathematical representation of a set of variables (which includes both models fit to data and formal models e.g. Robinaugh, Haslbeck, et al. (2019)), while a system refers to the combination of the data model and definition of control inputs. In our case, the aforementioned VARI model is being used to represent the dynamics governing the set of symptoms, and the system we present

below is an elaboration of that model which includes control inputs.

One of the simplest types of system is the discrete linear time-invariant system (DLTI). These systems operate over discrete time units, have no non-linear components (and can thus be represented entirely using linear algebra), and the system dynamics do not change over time (hence time-invariant). This type is well-suited for both representing network psychometric models (such as the a VARI model depicted in Fig. 1) and for illustrating control theory techniques. As such, we focus on DLTI in this paper (for more information about system identification, including the rationale for why we chose the ML-VARI model as the system of interest in this paper, see SM).

DLTI systems take the following form:

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t \quad (1)$$

$$\mathbf{y}_t = \mathbf{C}\mathbf{x}_t \quad (2)$$

where $\mathbf{x}_t$ is a vector of symptom (**state**) values at time t , $\mathbf{A}$ is the coefficient matrix governing their dynamics, and $\mathbf{B}$ is a $p \times l$ matrix mapping the interventions (**control inputs**) $\mathbf{u}_t$ to the symptoms. $\mathbf{y}_t$ is the vector of the observed measurements of the symptoms (**observations**) at time t while $\mathbf{C}$ is a matrix that maps the true values of symptoms $\mathbf{x}_t$ to the observed measurements $\mathbf{y}_t$. If one specifies that $\mathbf{C} = \mathbf{I}$ ($\mathbf{I}$ being the identity matrix) then $\mathbf{y}_t = \mathbf{x}_t$.

Although Eq. 2 is very similar to a vector autoregressive model at lag-1 that currently dominates the within-subjects network literature in psychopathology (Bringmann et al., 2013; Robinaugh, Haslbeck, et al., 2019)), there are three important differences: (a) the inclusion of the control inputs $\mathbf{u}_t$, (b) the addition of an observation equation, and (c) the lack of a disturbance term (ε_t in Eq. 2). The inclusion of control inputs is a necessity, though it is important to note they are additive inputs. The observation equation will feel familiar to readers versed in latent variable modeling, as it maps indicators (the observed measurements) to underlying states. To simplify our exposition, we make the assumption of complete observability throughout this manuscript. This assumption implies that there are no latent variables and that the observed variables are perfectly reliable, which is a problematic assumption when considering psychological constructs. Fortunately, more general forms of the $\mathbf{C}$ matrix can be specified to allow for latent $\mathbf{x}$ with manifest indicators $\mathbf{y}$. Using more general forms of $\mathbf{C}$ requires testing for *observability*, or the ability to reconstruct the latent variables $\mathbf{x}$ from the measured indicators $\mathbf{y}$. For a more detailed discussion of observability, see Preumont (1997).

Finally, in the original manufacturing context of control theory the absence of a disturbance term is defensible as disturbances are rare. In psychological systems, regular random disturbances are to be expected. Fortunately, however, if the random disturbance has an expected value of 0 and is uncorrelated in time (i.e., white noise), then disturbances can be ignored when considering optimal control inputs. The assumption of a white noise disturbance term is common across many time series and longitudinal models used in psychology and crucially relies on an assumption that any temporal dependence is explicitly modeled. Relaxing this assumption and allowing for time-correlated disturbances is the purview of *stochastic control theory* (See Åström (2006) for an introduction), a topic of relevance to psychological research, but not one we will

address here. To demonstrate the robustness of these methods to disturbances, we generate white noise disturbances for each of our simulations.

2. Using Control Theory to Select Intervention Targets.

With our system specified, we can leverage this information to identify which symptoms are the best intervention targets. For example, in our complicated grief network, it would be beneficial to know whether a targeted intervention to reduce, say, *emotional pain* is more effective than an intervention to increase *positive engagement*. However, the mere specification of the system is not sufficient to determine how an intervention will perform, as the performance of an intervention is entirely dependent on how it interacts with the dynamical system that it effects. Control theory provides tools to directly quantify the theoretical impacts of a targeted intervention in the form of *controllability centrality measures*. To begin, let us make some simplifying assumptions that reduce notational burden. First, let $\mathbf{C} = \mathbf{I}$. This reduces the system in Eq. 2 to

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t. \quad (3)$$

The matrix $\mathbf{B}$ describes how a set of interventions impacts the symptom values. From this system, we compute the controllability Gramian $\mathbf{W}_\infty$, which is the solution to the following Lyapunov equation in the case of a discrete time-invariant linear system (Lewis, Vrabie, & Syrmos, 2012):

$$\mathbf{W}_\infty - \mathbf{A}\mathbf{W}_\infty\mathbf{A}' = \mathbf{B}\mathbf{B}' \quad (4)$$

where $'$ is the transpose operator.

The controllability Gramian is related to efficacy of a given intervention with respect to its cost (*energy*). Eigenvectors with large eigenvalues define directions in the space of x that are less costly (e.g., time, money, effort, etc.) (Summers, Cortesi, & Lygeros, 2016). Interventions with larger values on the diagonal of $\mathbf{W}_\infty$ lead to more change in the overall system for an equivalent cost than do interventions with smaller values on the diagonal.

Critically, the controllability Gramian refers to the controllability of the entire system using the set of treatments specified in $\mathbf{B}$. As such, we can easily use the controllability Gramian to assess how efficiently the system can be controlled by specific symptoms. Let $\mathbf{B} = \mathbf{I}$. In this case, each intervention defined in $\mathbf{B}$ is the intervention on a single symptom. Therefore, the value on the diagonal of the controllability Gramian when $\mathbf{B} = \mathbf{I}$ quantifies the effect that intervening on each symptom individually has on the rest of the system (Summers et al., 2016). This controllability metric is called *average controllability*², which provides a direct measure of how efficient it is to target a given symptom for intervention, relative to the efficiency of targeting other symptoms in the symptom. An intervention that targets a high average controllability

²The use of the term “average” in average controllability is an unfortunate case of naming. Here, average does not refer to any specific statistical average, rather it is used in a more colloquial sense to mean “general” or “overall.” As such, there is the distinct possibility that researchers will use the term “mean average controllability” to refer to the average of several values of average controllability. We sympathize with the confusion of any reader who encounters that term in the future.

symptom will lead to more change in the system overall than an equivalently strong intervention that targets a symptom with lower controllability. Importantly, these are relative differences, and average controllability does not indicate that an intervention will have a strong or clinically meaningful effect. We elaborate on this point later in Section 3.1.

Average controllability is strongly associated with the network statistic of *strength* (the sum of the outgoing or incoming edge weights for a given node), which is commonly used when discussing intervention targets in psychological networks (Fried et al., 2017). However, average controllability holds three advantages over strength as a metric of symptom importance. First, whereas strength is concerned with network structure (specifically, the sum of connection strengths between a node and the rest of the network), average controllability is focused on the effect of a theoretical intervention, thereby more directly reflecting the ultimate end goal of identifying intervenable symptoms. Second, average controllability incorporates both the directionality of edges and the sign of edges into a single value per node. In contrast, strength must be separately calculated using either outgoing or incoming edges. Furthermore, most uses of strength do not account for the sign of edge weights (i.e., they take the absolute value). Failing to account for path signs can sometimes result in underestimating suppression effects, where a symptom that has high out-degree is not as effective to intervene upon as a symptom with lower out-degree, simply because the high out-degree symptom has a mixture of positive and negative effects. Finally, average controllability is *submodular* (Summers et al., 2016), meaning that, if one were interested in intervening on only k symptoms, then the best set of symptoms to intervene on are the k symptoms with the greatest average controllability. This guarantees that symptoms chosen on the basis of high average controllability should be as effective when intervened upon together as when they are intervened upon separately. In contrast, symptoms chosen on the basis of strength are more likely to be less effective intervention targets when intervened upon together.

Another useful controllability metric is *modal controllability* (Hamdan & Nayfeh, 1989). Whereas average controllability quantifies the efficiency of intervening on a symptom relative to intervening on other symptoms in a total symptom reduction sense, modal controllability quantifies the “reach” of intervening on a given symptom. The “modal” in modal controllability refers to the modes of the underlying dynamical system. Roughly speaking, modes refer to the independent directions a system can move in and are the dynamical systems analog to the orthogonal components of a principal components analysis in cross-sectional data.³ For example, an intervention that targets a symptom with low modal controllability would be more restricted in the ways it can change other symptoms relative to an intervention targeting high modal controllability symptoms. Because interventions targeting high modal controllability symptoms can effect the system in more ways, such interventions are also more likely to impact “hard-to-reach” symptoms than are less centrally located in the psychometric network.

To calculate modal controllability, let λ_j be the j th eigenvalue of $\mathbf{A}$ and ν_{ij} be the i th entry

³The term mode, when applied to a linear system, refers to the set of eigenvectors along with the corresponding eigenvalues of the dynamics matrix. These eigenvectors correspond to patterns of input that, if provided, would pass through the system unchanged save for a multiplicative constant, which is the corresponding eigenvalue. To be concrete, if ν is an eigenvector of A with corresponding eigenvalue λ , if $X_t = \nu$, then $X_{t+1} = \nu A = \lambda \nu$

on the j th eigenvector. The modal controllability of symptom i is then (Pasqualetti, Zampieri, & Bullo, 2014):

$$\phi_i = \sum_{j=1}^n (1 - |\lambda_j|^2) |\nu_{ij}|^2 \quad (5)$$

where n is the number of symptoms and $|\cdot|$ is the modulo operator for complex numbers. Modal controllability has a similar property to average controllability, in that one can simply rank order symptoms in terms of modal controllability to determine the set of k symptoms that maximize the smallest modal controllability (Pasqualetti et al., 2014).

Importantly, average and modal controllability can be used not only to assess interventions on single symptoms, but also broader intervention strategies that target multiple symptoms. This is especially useful in the context of predefined interventions, where the goal is not to determine which symptoms are best to target, but rather to determine which interventions (out of many) are most effective. Here again, we would stress that this is a significant advantage of controllability metrics over indices of node centrality, which only provide information about a single node in isolation and do not account for the effects of intervening on multiple symptoms, something that is almost certainly required in any real-world treatment context.

2.1. Generalizing Intervention Targets to the Population

To this point, we have seen the value of controllability metrics when working with both single- and multi-symptom interventions. However, we have not yet addressed a significant barrier in the treatment of psychological disorders. Psychological disorders exhibit considerable heterogeneity in both their symptom profile and system dynamics. A key question for any intervention design is: Are the symptoms best suited for intervention (i.e., most controllable) consistent across the population, or is there heterogeneity in the best treatment targets?

Within the framework of network psychometrics, this heterogeneity means that our dynamic network $\mathbf{A}$ is a random variable, and therefore any network statistic derived from it is also a random variable. Simply because a set of symptoms shows a particular ordering of average or modal controllability (computed from a single subject's within-subject network) does not imply that this ordering generalizes to another individual in the population. To evaluate if the ordering of controllability statistics is consistent in a sample of within-subject networks, further analyses are needed. We developed a simple testing procedure to do just this.

By testing rank order differences between all pairwise combinations of symptoms, followed by a false discovery rate correction, we can determine if the observed ranking of the mean centrality measures is generalizable across subjects, or if certain symptoms are indistinguishable by rank in the population. Let $\mathbf{A}_1, \dots, \mathbf{A}_n$ be a sample of within-subject networks derived from n subjects, and let x_p be a n length vector of a centrality measure computed for a given symptom p . Then, for every pairwise combination of symptoms, perform a Wilcoxon signed-rank sum test (Wilcoxon, 1945). When there are k symptoms in the networks, this leads to $k(k-1)/2$ dependent tests, which can be corrected for false discovery rate using the Benjamini-Hochberg correction (Benjamini & Hochberg, 1995). A significant Wilcoxon signed-rank sum test comparing symptom k to symptom l indicates that the rankings of symptoms k and l are significantly different.

FIGURE 2.

Tests of Symptom Ranking on Average and Modal Controllability. Symptoms are sorted left to right in ascending order of the controllability centrality value. Tile colors represent the p -value of the associated Wilcoxon signed-rank sum test (after FDR correction). Black: $p < .05$, Grey: $.05 \leq p \leq .1$, Red: $p > .1$. Average controllability (left panel) results show that *suicidal ideation* (SI) and *thoughts of the deceased loved one* are the symptoms with the highest average controllability, that they are not distinguishable in rank from one another, but are consistently higher in average controllability than any other symptom.

In Fig. 2, we show the results of this testing procedure for both average controllability and modal controllability applied to our within-subject complex grief symptom networks estimated with the ML-VARI approach (see Fig. 1). Recall that controllability is not how well a symptom can be controlled, but rather how controlling a given symptom might impact the rest of the system. Our results show that for average controllability, the symptom ranked highest is *suicidal ideation*. However, this symptom is not significantly different from the next highest symptom, *thoughts: loved one* (i.e., thoughts about the deceased loved one), suggesting that there is no clear symptom that would be the single most effective for intervention. Rather, based on these findings, an intervention targeting both symptoms would likely prove most effective across the population.

These analyses also suggest that *thoughts: future* (thoughts about one’s own future), *thoughts: death*, *yearning: loved one*, and *positive activities* (engagement in positive activities, reverse coded) are the next set of highest average controllability symptoms. This is not a clean grouping, with many non-significant rank differences between these symptoms and lower rank symptoms. Nonetheless, the existing grouping suggests that if one could not intervene on *suicidal ideation* and/or *thoughts: loved one* directly, then intervening on these four symptoms would be the next best approach. This highlights the system-level concerns of controllability, as the relations between all symptoms are relevant when considering which symptom to intervene on.

In contrast to average controllability, the ordering of modal controllability (right panel of Fig. 2) was less clear, likely because modal controllability is more sensitive to minor differences in within-subject network structure than average controllability. One clear finding is that *suicidal ideation* has the highest modal controllability rank, and this ranking is significantly different from

the lower half of symptoms. Additionally, *thoughts: loved one* is again ranked second with a similar pattern of non-significant differences, further suggesting that an intervention targeting both *suicidal ideation* and *thoughts: loved one* will be more effective at changing the system in this group of participants relative to alternative strategies that target any two other symptoms.

3. Using Control Theory to Simulate Treatments

As the example of the complicated grief network illustrates, controllability measures provide a powerful tool for selecting intervention targets. By selecting symptoms with consistently high controllability, intervention developers can maximize efficient use of resources. However, the identification of treatment targets is only one advantage of a control theory approach. A second important advantage is that control theory provides the opportunity to evaluate how well interventions will perform, even before empirical studies such as randomized controlled trials are carried out. We can achieve this aim using *optimal control methods* that seek to determine a sequence of control inputs that will drive a system to a desired end state with minimum cost. To introduce these methods, we will review three key concepts: tests of controllability, cost functions, and optimal control.

3.1. Testing Controllability

One fundamental question is whether a given system is controllable. Is it possible to control the trajectory of all the symptoms over time, given a set of interventions and the dynamics of the system? Note that *controllability* does not refer to the ability to control a single symptom. Rather, it refers to the ability to control the entire system from a single symptom or set of symptoms. A system is *controllable* if given $\mathbf{A}$ and $\mathbf{B}$, for any desired final state $\hat{\mathbf{x}}$, there exists a series of interventions (*control inputs*) such that $\mathbf{x}_{\hat{t}} = \hat{\mathbf{x}}$ for some $\hat{t} > 0$. Although this property is important to evaluate, it doesn't reflect the practicality of controlling the system. A system can be mathematically controllable, but require a unreasonable amount of time or intervention strength. Conversely, a system might not be mathematically controllable, but a given intervention still proves to be effective at reducing symptoms by a clinically relevant degree. Evaluating the controllability of a system provides useful diagnostic measures, but further steps are necessary to evaluate the actual effect of a given intervention. There are two approaches to assessing controllability that we present here, the first being a classical test of controllability and the second being a descriptive approach using indirect effect calculations.

To test if a system is controllable, one can construct the following matrix:

$$\mathcal{C} = [\mathbf{B}, \mathbf{AB}, \mathbf{A}^2\mathbf{B}, \dots, \mathbf{A}^{n-1}\mathbf{B}] \quad (6)$$

where n is the number of variables in the system. The system is controllable if $\mathcal{C}$ has full row-rank $\text{rank}(\mathcal{C}) = n$. When we consider a single intervention that reduces the symptom of complicated grief with the highest average controllability (i.e., *suicidal ideation*) equally, we find that the system is not controllable, as $\mathcal{C}$ is not full rank. This implies that the intervention targeting *suicidal ideation* is not capable of reducing all symptoms to 0, no matter what sequence of intervention strengths are applied.

For an alternative and perhaps more intuitive approach to examining controllability, we can take advantage of the fact that $\mathbf{A}^i \mathbf{B}$ represents the effect of the interventions described in $\mathbf{B}$ after i timesteps. Therefore

$$\mathbf{E}_\infty = \mathbf{B} + \sum_{i=1}^{\infty} \mathbf{A}^i \mathbf{B} \quad (7)$$

is a $n \times d$ matrix where entry $\mathbf{E}_{\infty[ij]}$ represents the total effect of intervention j on symptom i if the intervention is applied once at time $i = 0$. This matrix provides a practical look at the overall effect of an intervention and can indicate whether an intervention will be effective at changing a given symptom.

Considering the same intervention that targets *suicidal ideation*, we can calculate the intervention effect matrix $\mathbf{E}_\infty$, which is presented in Table 2.

TABLE 2.

Symptom	Suicidal Ideation	Total Potential Intervention Effect								
		Thoughts: Loved One	Yearning: Loved One	Approach	Thoughts: Death	Emotional Pain	Avoidance	Thoughts: Future	Pleasure	Positive Activities
Treatment Effect	-2.09	-0.63	0.10	-0.05	0.138	0.123	-0.01	0.08	0.13	0.15

Table 2 indicates that a single intervention that reduces *suicidal ideation* by 1 point, will lead to a total reduction over time of 2.22 points for *suicidal ideation* and 2.36 points for *thoughts: loved ones*. However, the same intervention would lead to, for example, an *increase* in *yearning: loved ones* of 0.24. As there is a single intervention and that intervention reduces some symptoms while increasing others, it does not appear that this system (i.e., complicated grief) can be wholly controlled by a single intervention that targets suicidal ideation.

If an single-node intervention *suicidal ideation* failed to achieve control of the system. The next question is: What is the minimal set of interventions that is sufficient to control the system? Unfortunately, there is no simple solution to this problem. Determining this set exactly (**actuator placement**) is computationally intractable (indeed, an NP-hard problem, Olshevsky, 2014a), though there are algorithms that find a set of control inputs that render the system controllable (Olshevsky, 2014a, 2014b). There are practical barriers as well. First, it might be impossible to directly intervene on some symptoms. For example, in Cognitive Behavioral Therapy, interventions on emotions are only targeted indirectly via interventions on cognition or behavior. Second, it is unlikely that any given intervention will only directly causally impact a specific symptom or subset of symptoms, with other symptoms being impacted indirectly. Instead, interventions likely directly effect many symptoms within a system to varying degrees of strength. We illustrate the impact of assuming a highly specific intervention vs one with more broad effects in Section 4.1 and show how the treatment effects of a real world intervention suggest non-specificity in Section 4.4.

3.2. The Cost Function

Until now the tools we have described for evaluating the effect of an intervention, controllability centrality measures and tests of mathematical controllability, do not account for

the cost of an intervention. Both approaches operate under the assumption that an intervention can be arbitrarily applied at any strength. These tools are still useful of course, as they provide evaluations of the potential efficacy of an intervention, but do not provide any guidance as to the implementation of a given intervention. In order to apply control in any practical context, we first need to define how costly our control inputs are. To do this, we use a *cost function*. A general form of the cost function that provides analytic solutions is that of a linear-quadratic regulator (Lewis et al., 2012):

$$J_t = \frac{1}{2} \mathbf{x}'_T \mathbf{S}_T \mathbf{x}_T + \frac{1}{2} \sum_{k=t}^{T-1} \mathbf{x}'_k \mathbf{Q}_k \mathbf{x}_k + \mathbf{u}'_k \mathbf{R}_k \mathbf{u}_k \quad (8)$$

where T is the final timepoint. Many cost functions in the context of control theory are similar to the general linear-quadratic regulator, as this form is particularly amenable to analytic solutions to the control problem. In practice, choice of the cost function depends on the nature of the control problem. Here, we are seeking to minimize symptom values.

In the above cost function, the matrices $\mathbf{S}_T$, $\mathbf{S}_k$ and $\mathbf{R}_k$ have important substantive interpretations. $\mathbf{S}_T$ is a $p \times p$ positive (semi) definite matrix that provides the relative weighting of the final values of each symptom. For example, $\mathbf{S}_T = \mathbf{I}$ places equal importance on each symptom. The sequence of matrices $\mathbf{Q}_k$ are all $p \times p$ positive (semi) definite matrices specifying the relative importance of each symptom at each timepoint. $\mathbf{Q}_k$ can be specified as a single time-invariant matrix or set to 0, in which case the intermediate values of the symptoms is not taken into account when minimizing the cost function. Finally $\mathbf{R}_k$ is a sequence of $p_u \times p_u$ positive (semi) definite matrices specifying the relative cost of each intervention. As with $\mathbf{Q}_k$ this can be specified as time-invariant but cannot be set to 0. By manipulating $\mathbf{S}_T$, $\mathbf{Q}_k$ and $\mathbf{R}_k$, researchers can examine the theoretical performance of their intervention set under a variety of settings. To simplify notation, $\mathbf{Q}$ and $\mathbf{R}$ will be considered as time-invariant.

Many cost functions in the context of control theory are similar to the general linear-quadratic regulator, as this form is particularly amenable to analytic solutions to the control problem. In practice, choice of the cost function depends on the nature of the control problem. Here, we are seeking to minimize symptom values both by the final timepoint ($\mathbf{S}_T$) and during the period of control ($\mathbf{Q}_t$), while keeping cost of the intervention minimal ($\mathbf{R}$). As such, this function is a reasonable choice for the control problem at hand.

3.2.1. Tuning Cost Parameters

Unlike parameters that govern dynamics (i.e., $\mathbf{A}$ and $\mathbf{B}$), which could be inferred as part of system identification, the parameters of the cost function J must be defined by the analyst. There are several considerations for this definition:

- **Relative Weighting** - Each of the matrices $\mathbf{S}_T$, $\mathbf{Q}_t$ and $\mathbf{R}$ contribute to the total cost, and when specifying them, the analyst must consider the relative weighting both within and between matrices carefully. For example, if $\mathbf{S}_T$ and $\mathbf{Q}_t$ are $\mathbf{I}$ while $\mathbf{R} = 10\mathbf{I}$, then the cost of applying interventions will be weighted more heavily than the values of symptoms.
- **Consider Scaling** - If symptoms are measured on a different scale (i.e., one symptom is on a 1-7 point scale, the next a 1-5 point scale), this will affect how the control parameters

function. For example, weighting the importance in the $\mathbf{S}_T$ matrix of 1-5 symptom and a 1-7 symptom equally (e.g., diagonal entries of $\mathbf{S}_T$ equal 1 for those symptoms) will result in the symptoms having **unequal importance**, as the scaling is different.

- **Consider Timeframe** - The relative weights of $\mathbf{R}$ vs. $\mathbf{S}_T$ and $\mathbf{Q}_t$ will affect the timescale of the intervention strategy. If $\mathbf{R}$ makes much more of a contribution to cost than symptoms do, then an "optimal" intervention strategy might be to apply the intervention at very low strength, which requires a longer intervention period. Alternatively, if cost of interventions is not weighted strongly, then the intervention might be applied too strongly or too quickly.

There are no clear cut rules for setting the control parameter values. The precise values that an analyst chooses will be dependent on the data and purpose at hand. We emphasize that the choice of values is not an inferential process, as they do not depend on the data generating process but rather reflect the desired goals of applying control.

3.3. Optimal Control

With our cost function J defined and parameter values set, we can now determine the optimal series of control inputs that moves the system to the target state while minimizing the cost function in a finite time T . We do not present derivations here, but for a thorough treatment, see Lewis et al. (2012). First, we must compute an intermediate series of weighting matrices:

$$\mathbf{S}_t = \mathbf{A}'(\mathbf{S}_{t+1}^{-1} + \mathbf{B}\mathbf{R}^{-1}\mathbf{B}')^{-1}\mathbf{A} + \mathbf{Q}. \quad (9)$$

Note that $\mathbf{S}_t$ is defined with recursion forward in time, where the value of $\mathbf{S}_t$ at time T is the *a priori* defined $\mathbf{S}_T$. The above expression for $\mathbf{S}_t$ is referred to as a Riccati equation in the control theory literature.

With $\mathbf{S}_t$ computed for every $t \in 1 \dots T$, we can then compute a set of matrices that allow us to inform our intervention strategy based on the current symptom values (**Kalman gain sequence**):

$$\mathbf{K}_t = (\mathbf{B}'\mathbf{S}_{t+1}\mathbf{B} + \mathbf{R})^{-1}\mathbf{B}'\mathbf{S}_{t+1}\mathbf{A} \quad (10)$$

and finally the optimal set of interventions

$$\mathbf{u}_t = -\mathbf{K}_t\mathbf{x}_t. \quad (11)$$

Note that the optimal sequence of control inputs are computed on the basis of the observed sequence of $\mathbf{x}$. In other words, the optimal sequence of interventions is a closed feedback loop, as the values of $\mathbf{x}_t$ inform the control inputs $\mathbf{u}_t$. Importantly, this differs from open loop controls, which can be used to drive systems to predefined endpoints instead of minimizing the cost function (i.e., by time T , all symptoms will be 0). However, because open loop controls do not take into account the current value of the symptoms, they are not robust to any deviation off of the expected trajectory, making them unsuitable for controlling an inherently noisy and measurement error prone system like that of a psychological disorder.

3.4. Two Simple Control Examples

To demonstrate the use optimal control methods, we will first apply them in two simple simulated examples: an easy-to-control (System 1) and a difficult-to-control (System 2) system. Both systems have three variables (A, B, and C), a single intervention that is applied to variable A, and a series of white noise disturbances (randomly generated once and applied to both System 1 and System 2). $\mathbf{S}_T$ and $\mathbf{Q}_t$ are both $\mathbf{I}$, and $\mathbf{R} = 100$. The VARI networks for both systems, the control inputs, and the variable time series appear in Figure 3. In both cases, the starting values for all three variables was set at 20. The key differences between System 1 and System 2 are the presence of a suppressing effect from C to B and the negative autoregressive effect on B, both of which render System 2 more difficult to control than System 1.

These simulations illustrate two important points. First, this form of closed-loop feedback control is fairly robust to disturbances. In both systems, even with disturbances, control successfully minimized variable A and further reduced the value of variable C, keeping both fairly low across time. Second, a system can be controllable in an absolute sense but not controllable in a practical sense. In the difficult-to-control System 2, our interventions had minimal effect. Even though variable A was positively associated with variable B in System 2, the negative relation between C and B partially suppresses the A-to-C relation, and the negative autoregressive effect of B further dampens any interventions that impact B through A. This is the case even though both System 1 and 2 are “controllable,” as indicated by the previously described rank test for controllability. Examining the total intervention strength for both systems, an intervention that reduces variable A by 1 would lead to a total reduction in variable B of 2.88 in System 1 but only .1 in System 2, an order of magnitude different. In other words, even though System 2 could be controlled with very high cost and/or over a very long timeframe, it is not controllable in a practical sense.

FIGURE 3.

Networks, Intervention and Variable Values for Systems 1 and 2. Networks represent dynamics operating on change (as in Figure 1). Both systems were subjected to the same series of disturbances, which was generated as white noise.

4. Illustrating the Clinical Utility of Control Theory

4.1. Controlling Psychological Networks via Single-Target Intervention

Now that we have established that closed-loop interventions work in principle in simple simulated systems, we turn back to the VARI model of complicated grief symptoms estimated from empirical data. Although the network psychometric modeling literature has often focused on identifying symptoms thought to have the largest potential impact on other symptoms, another potential goal of these models could be to identify the specific symptoms most amenable to intervention. The simplest scenario is an intervention that targets a single symptom. Previously, we identified *suicidal ideation* as the symptom with the highest average controllability (indicating that *suicidal ideation* is the most cost efficient symptom to control). In Section 3.1 we showed that a *suicidal ideation*-only intervention was not capable of controlling the entire system. Here, we demonstrate what that means in practice.

To begin setting up the matrices that form the cost function for this simulation, we assume that (a) no symptom is more important to reduce than any other and (b) this importance does not change over the course of treatment. This implies that $\mathbf{S}_T = \mathbf{Q} = \mathbf{I}^4$ and when $\mathbf{R} = a$, where a is a positive number chosen to ensure that intervention is not unrealistically strong. The control input matrix $\mathbf{B}$ is, for a single symptom intervention targeting *suicidal ideation*, a 10×1 vector of 0s, with the exception of the element corresponding to *suicidal ideation*, which is set to 1. Our dynamics matrix A is taken from the group network of complicated grief symptoms (Fig. 1), see SMs for more details. We add in a small amount of random disturbance to each symptom at each timepoint, sampled independently from a $Normal(0, .1)$ distribution. In the following simulations, all symptoms will start at their middle value (3 for *suicidal ideation* and 5 for every other symptom), and are restricted in their ranges from 0 to the symptom maximum (5 for *suicidal ideation*, 7 for all other symptoms). Finally, intervention strength (**control inputs**) are restricted to be positive or zero. This restriction prevents the case where the optimal control would involve increasing symptoms by applying an "anti-intervention."

The control input matrix $\mathbf{B}$ is, for a single symptom intervention targeting *suicidal ideation*, a 10×1 vector of 0s, with the exception of the element corresponding to *suicidal ideation*, which is set to 1. Previously, we raised the point that a psychological intervention, even one explicitly designed to reduce a single symptom, is unlikely to have *no* direct effects on any other symptoms. To reflect this, we can modify the *suicidal ideation*-targeted intervention to have it reduce all other symptoms by a fraction of what it reduces *suicidal ideation*. In Figure 4, we show what optimal control looks like when the treatment reduces *suicidal ideation* by 1 point and all other symptoms by .1 points, while in Figure 5 we show what optimal control looks like when the intervention only reduces suicidal ideation.

⁴It is vital that these matrices are specified with respect to the scaling of individual symptom measures. In the example of complicated grief symptoms, the measure of *suicidal ideation* has a maximum value of 5, whereas all others measures have a maximum value of 7. This requires a slight modification of the $\mathbf{S}_T$ and $\mathbf{Q}$ matrices, so that the scales are equated. Here, the element corresponding to *suicidal ideation* in $\mathbf{S}_T$ and $\mathbf{Q}$ are set to $\frac{7}{5} = 1.4$ rather than 1. For scales with several different maximums, a common factor can be computed. Finally, this rescaling should only be done when the choice of units is arbitrary, such as in a Likert type item

First, let us examine the symptoms' behavior and the intervention's magnitude over the course of 100 timepoints. $\mathbf{R}$ is set to 100, which (in this example) restricted the maximum intervention strength to an approximately .3 point reduction of *suicidal ideation*.

FIGURE 4.
Resulting symptom values and intervention strength when intervening on suicidal ideation with strength of 1, and all other symptoms with strength of .1.

FIGURE 5.

Resulting symptom values and intervention strength when only intervening on suicidal ideation. Note that the majority of the reduction in *suicidal ideation* occurs within the first 15 timepoints, and that symptoms other than *suicidal ideation* and *thoughts: loved one* are not significantly effected by the *suicidal ideation* only intervention.

Figure 4 shows that the intervention reduces all symptoms by a small amount. There are three main differences between Figures 5 (intervention directly reduces only suicidal ideation) and Figure 4 (control additionally slightly reduces all other symptoms). First, the intervention is applied at low strength continuously throughout the timeframe. Second, *thoughts: loved ones* showed a greater reduction and less variance, as did *suicidal ideation*. Finally, the remaining symptoms had a slightly smaller spread in their trajectories, though the majority of the remaining symptoms did not show meaningful reductions.

Although *suicidal ideation* was the symptom with the highest average controllability for both

within- and between-subject networks, the symptom trajectories in Fig. 5 and Fig. 4 show that an intervention that solely targets *suicidal ideation* reduces *suicidal ideation* and *thoughts: loved ones*, but no other symptoms, and that this occurs even when the intervention directly reduces all other symptoms by a small amount. These results highlight two crucial points. First, controllability metrics, such as average controllability, indicate how efficiently interventions targeting single symptoms will impact the whole system *relative to targeting other symptoms*. Second, although controllability metrics can help us identify the most impactful symptoms, they indicate little about the actual impacts we should expect from such an intervention. Selecting intervention targets is important, but testing these interventions is necessary to gain a genuine understanding of the impact they will have. This example also illustrates how control theory can guide the application of maintenance interventions. For example, the controller applied the intervention around timepoint 75 after a period of relatively low intervention strength. This was in response to the slight trend upwards seen in *suicidal ideation* and *thoughts: loved ones*.

To show that an intervention on *suicidal ideation* is more effective than other single symptom interventions we can compute an *intervention efficacy ratio* (IER):

$$\text{IER} = \frac{\text{Total Reduction in Symptoms}}{\text{Total Intervention Strength}} \quad (12)$$

for each single target intervention, while setting the cost ($\mathbf{R}$) to be equal, and comparing these ratios across symptoms. To calculate this, we re-simulated the trajectories in Figure 5 without any disturbance component, and calculated out the total reduction in all symptoms from the initial timepoint to the final timepoint, and the sum of the intervention strength across the entire timeframe. This ratio is a direct measure of the theoretical efficacy for a given intervention with units of reduction in symptoms per unit of intervention strength, at a given $\mathbf{S}_T$, $\mathbf{Q}$ and $\mathbf{R}$. The IER for a *suicidal ideation* intervention was 1.75, the IER for a *thoughts: loved one* intervention is .87, and all other single symptom interventions have IERs close to .3. This in turn indicates that although a *suicidal ideation*-targeted intervention was not capable of reducing all symptoms to 0, *it was still the most effective intervention in a minimal cost, maximum change sense*. Note that IERs can be calculated for any intervention, including ones that target multiple symptoms.

Overall, our simulations and IER findings suggest that applying single target interventions can be effective in reducing the targeted symptom, but are less likely to meaningfully impact all symptoms, even when the intervention has small direct effects on non-targeted symptoms. However, it is unlikely that a single intervention is all that a therapist has on hand. Instead, there are likely several available interventions, some of which can be applied simultaneously or in sequence, a situation we cover next.

4.2. Controlling Networks via Multiple-Target Interventions

Suppose Intervention A targets *suicidal ideation*, Intervention B targets *positive activities* (reverse coded so that reductions correspond to increases in enjoyment of positive activities), and Intervention C targets *thoughts: future*. In this example, all three interventions also reduce non-targeted symptoms by .1. The $\mathbf{B}$ matrix for these interventions is shown in Table 3.

TABLE 3.
B Matrix for Interventions A, B and C.

	Intervention		
	A	B	C
Suicidal Ideation	-.1	-.1	-.1
Thoughts: Loved One	-.1	-.1	-.1
Thoughts: Future	-.1	-.1	-.1
Avoidance	-.1	-.1	-.1
Thoughts: Death	-.1	-.1	-.1
Positive Activities	-.1	-.1	-.1
Yearning: Loved One	-.1	-.1	-.1
Approach	-.1	-.1	-.1
Pleasure	-.1	-.1	-.1
Emotional Pain	-.1	-.1	-.1

We apply these three interventions simultaneously, with all symptoms being weighted equally, and all interventions having the same cost (e.g. $\mathbf{R} = 100\mathbf{I}$). Figure 6 shows the intervention strength and symptom trajectories in this case.

FIGURE 6.

Symptom Values and Intervention Strength for Interventions A, B and C. Intervention A targets *suicidal ideation*. Intervention B targets *positive activities*. Intervention C targets *thoughts: future*.

First, Figure 6 indicates that when these three interventions are applied simultaneously, all symptoms can be meaningfully reduced within the given timeframe. Recall that this was not the case for any of these three interventions when applied alone, as discussed in the previous subsection. Second, the interventions are not applied equally strongly. Intervention A is applied with the weakest strength throughout the timeframe, while Interventions B and C are applied with equal strength throughout the timeframe. Third, this set of interventions only reduced *suicidal ideation*, *thoughts: loved ones*, *thoughts: future* and *avoidance* to near 0. All other symptoms were systematically reduced, and more importantly, stabilized at that reduced level. This is notable as *positive activities* was the specific target of Intervention B, yet was not successfully reduced to 0.

where for both $\mathbf{S}_T$ and $\mathbf{Q}$, the first column/row corresponds to *suicidal ideation*.

$$\mathbf{R} = a \begin{bmatrix} 2 & 0 & 2 \\ 0 & 1 & 0 \\ 0 & 0 & .5 \end{bmatrix} \quad (15)$$

where the rows/columns are ordered, A, B, C. This $\mathbf{R}$ matrix puts our assumptions about relative cost of interventions into rigorous form. The 1 corresponds to B being half as expensive as A, while the .5 corresponds to C being a quarter as expensive as A. To keep intervention strengths from being unrealistically high, $a = 2500$.

Fig. 7 shows the symptom trajectories and intervention strengths for this control setup. Changing the control parameters led to these interventions being applied quite differently than in Figure 6. First, the fact that Intervention A and C are difficult to apply together is reflected in how these two interventions were applied. There is a spike in the strength of Intervention C around time 10, and there is a spike in the strength of Intervention A around time 65. Both interventions were applied at a very low level throughout the timeframe as well. Second, even though SI was set as the most important symptom to reduce, the inability to apply all three interventions simultaneously resulted in SI being reduced at a slower rate than that shown in 6. Finally, while the same sets of symptoms as in 6 were reduced to near 0, the remaining symptoms were not reduced as much overall as in 7.

In these examples we have been evaluating interventions that have been proposed on the basis of controllability statistics. However, it is common that we have a priori knowledge on the efficacy of real-world interventions. In the next section, we simulate the effects of the complicated grief treatment intervention Shear (2010) that was given to our sample of patients.

FIGURE 7.

Symptom Values and Intervention Strength for Interventions A, B, C with Unequal Symptom Importance and Intervention Costs. Note: Interventions A and C have been specified to be incompatible (i.e. applying them simultaneously costs more than applying both separately)

4.4. Evaluating Empirical Interventions: A Treatment for Complicated Grief

The control theory methods outlined herein allow intervention developers to propose a theoretical intervention and evaluate its efficacy within a simulation framework. Such simulations can then guide the design of a randomized control trial, from which empirical data regarding intervention performance are collected. Ideally, researchers will then use that empirical performance to inform a follow-up control theory analysis (essentially, cross-validation) of that intervention's potential efficacy. This allows researchers to formalize predictions of treatment

response, evaluate how well those predictions match treatment response, and improve future predictions for other intervention studies. By performing this cross-validation of the intervention effect, researchers can better ensure their intervention's effects generalize to the population.

In our complicated grief example, the 8 patients with complicated grief and persistent suicidal thoughts received a 16-session complicated grief treatment intervention (Shear, 2010). Throughout this treatment, patients completed a range of treatment components, including psychoeducation about grief and an overview of complicated grief treatment (Sessions 1-3); imaginal and situational revisiting exercises in which they confront and process the reality of their loss and reduce grief-related avoidance (Sessions 4-10); and develop plans for the future, in terms of both meaningful, positive long-term goals and plans for managing times in which grief can be especially difficult (Sessions 11-16). For further discussion of the treatment, its rationale, and components, see (Shear, 2010).

To estimate the effect of the applied treatment, a treatment indicator was added as a fixed exogenous predictor of each symptom in the original VARI model. Although this is a simple way of estimating a treatment effect, it does assume that the effect of the intervention is the same for all participants. For a discussion about the difficulties and advantages of using individualized treatment effect estimates, see our Discussion section. Table 4 shows the symptom-level effects this treatment on these 8 patients, with the largest decrease in *thoughts: loved one* (for details, see SMs).

For comparison, Fig. 8 shows the result when we apply control theory on the group-level complicated grief network estimated previously (Fig. 1). In this case, $\mathbf{R}$ is set to 150 to set the maximum intervention strength to be below 1. No symptom was set to be more important than any others (i.e. $\mathbf{S} = \mathbf{Q} = \mathbf{I}$). The results suggest that the empirical treatment is most effective at reducing *thoughts: loved one* (in terms of total reduction), as well as reducing *suicidal ideation* to near 0. Several other symptoms were reduced slightly. Finally, the IER of the empirical treatment (which recall is an estimate of total reduction in symptoms for a 1 unit application of the intervention) is .81, which compares favorably to the expected direct reduction in symptoms (.56; sum of the treatment effects in Table 4).

TABLE 4.
Intervention Effects for Empirical Treatment

Symptom	Suicidal Ideation	Thoughts: Loved One	Yearning: Loved One	Approach	Thoughts: Death	Emotional Pain	Avoidance	Thoughts: Future	Pleasure	Positive Activities
Treatment Effect	-0.057	-0.204	-0.024	-0.120	-0.089	-0.001	-0.038	-0.009	-0.071	0.053

FIGURE 8.
Symptom Values and Intervention Strength for Empirical Intervention.

4.5. Evaluating Between-Subject Heterogeneity in Treatment Response

One key assumption of the above simulations is that the dynamics matrix $\mathbf{A}$ estimated at the whole sample level is a reasonable representation of any given subject's dynamics. That is, we assume that there is little heterogeneity across different patients in the temporal dynamics among symptoms. However, mental disorders often exhibit pronounced inter-individual differences (Fried & Nesse, 2015; Young, Lareau, & Pierre, 2014; Zimmerman, Ellison, Young, Chelminski, & Dalrymple, 2015), including in their temporal dynamics (Fisher, Medaglia, & Jeronimus, 2018; Fisher, Reeves, Lawyer, Medaglia, & Rubel, 2017). Although there is little empirical work on heterogeneity for complicated grief, the sheer number of unique symptom combinations one can

have and still meet diagnostic criteria for the disorder (37,650; Lenferink & Eisma, 2018) suggests it is unlikely to be an exception. This means that modeling and evaluating inter-individual differences in treatment response is vital for understanding how effective an intervention will be when applied across many patients. Control simulations using within-subject dynamic matrices allow us to address this concern.

In our complicated grief example, we can estimate one dynamic network for each of the 8 participants rather than for the entire group (Fisher et al., 2017; Frumkin, Piccirillo, Beck, Grossman, & Rodebaugh, 2019). For calculation details, see SMs. Here, we wish to assess response heterogeneity for the empirical treatment described above. As *suicidal ideation* is of particular treatment importance for complicated grief, we focus once again on how the applied intervention impacts *suicidal ideation*.

Fig. 9 shows the intervention strength and reduction in *suicidal ideation* for all 8 subjects. These simulations are based on the actual individual network structures of each subject, but assume for simplification and to facilitate comparability, that all 8 subjects start at the same values (3 for *suicidal ideation*, 5 for all other symptoms) and that they experience the same set of disturbances over time ⁵. The general pattern of intervention strength was similar for all subjects, although the rate of *suicidal ideation* decreased and the volatility of *suicidal ideation*'s movement through time differed considerably between subjects.

For example, *suicidal ideation* for ID-6 was effectively reduced to 0, which indicates that the intervention would likely be effective for ID-6. However, ID-1 has rather minimal reductions in *suicidal ideation* over the timecourse. This ability to probe the expected effects of an intervention at the idiographic level informs efforts to personalize treatments in two ways. First, it allows researchers to quantify heterogeneity in treatment response when designing new intervention strategies (e.g., prioritizing the dissemination of an intervention with moderate but consistent effects over an intervention that works well for some but poorly for many). Second, it equips researchers to move toward the aim of personalized treatment, identifying the intervention strategy best suited for any given individual.

⁵Each subject experiencing the same set of disturbances is unrealistic for a real world intervention setting and was used in this simulation to emphasize that the intraindividual differences in system dynamics matter even in the presence of disturbances. If disturbances are generated independently for each subject, the snowballing impact of those disturbances can mask relevant differences in individual treatment response.

FIGURE 9.

Intervention strength and expected ‘Suicidal Ideation’ trajectories for 8 within-subject system dynamics.
 Note: the sequence of random disturbances applied to symptom trajectories was the same for all subjects.

5. Discussion

The use of network psychometric models to inform intervention design is a promising aim for clinical research, but one that requires new mathematical tools if it is to be realized. In this paper, we introduced several fundamental concepts from control theory, such as controllability centrality and optimal control, and illustrated via examples how they can improve intervention evaluation and design. All examples used EMA data from a sample of subjects undergoing treatment for complicated grief; this allowed us to firmly embed the methodology discussed in a practical, real-world empirical example. Below, we summarize our substantive findings and propose future

directions for methodological research in control theory and network psychometrics.

5.1. Controlling Psychological Networks: The Empirical Case of Complicated Grief

Our analyses, although meant in part to demonstrate methodology, revealed several interesting features of complicated grief and its treatment. Suicidal ideation and thoughts about one's deceased loved one had the highest average controllability, suggesting that intervening on those two symptoms in the present sample would be most effective at changing the system of complicated grief as a whole. However, an analysis of single symptom interventions showed that although targeting suicidal ideation was indeed the most effective single symptom intervention, it was not successful at reducing all other symptoms to zero. This suggests that the dynamic system that underlies complicated grief is not solely driven by any one symptom, and that single symptom interventions would not be effective at treating the disorder's gestalt.

Modeling and comparing multiple interventions sets (A, B, C) shed further light on the structure of complicated grief. In the first set of interventions, Intervention A (targeting the symptoms of *suicidal ideation* and *thoughts: loved one*) was more effective, in terms of intervention cost to total reduction in symptoms, relative to Interventions B or C when applied separately. However, when applied in tandem, Interventions A, B and C together led to reductions in all symptoms. Finally, the treatment that the subjects were undergoing—a structured psychological therapy focused on accepting the loved one's death—appeared to primarily impact patients' preoccupation with thoughts of the deceased loved one while also significantly reducing suicidal ideation in the process. All models were computed under the assumption that the group-level dynamics were a reasonable estimate for any subject in the population.

When accounting for inter-individual differences in temporal dynamics, we found that a wide range of intervention sensitivity in the present sample. Under the simulated intervention based on the empirically estimated treatment effects, several subjects did not exhibit much reduction in suicidal ideation, while others improved substantially. This highlights the need to model and ultimately understand psychological disorders at the level of the individual when designing interventions.

These substantive conclusions should be interpreted in light of a number of limitations. First, there are only eight subjects present in the analysis, each of whom reported not only complicated grief, but also suicidal thoughts, limiting the generalizability of our findings. Next, the treatment effect was modeled within the ML-VARI as a 0/1 indicator, which does not reflect the fact that the treatment had multiple phases. This choice, made to simplify the control theory modeling of the treatment effect, likely also reduces our power to detect treatment effects that varied over the stages of the intervention. Additionally, although we used an integrated VAR(1,1) model to model symptom trajectories as a non-stationary process, it is unlikely that this model truly represents the causal dynamics underlying the disorder, a significant limitation to which we will return.

5.2. Methodological Challenges and Future Directions

In this paper, we aimed to introduce control theory as a valuable tool for identifying intervention targets, understanding the behavior of a psychological system undergoing

intervention, and evaluating heterogeneity of intervention response. Below, we more fully discuss limitations in the presented approach, briefly introduce more modern control theory methods, and identify future research directions that may radically improve the application of control theory to psychology.

5.2.1. The Importance of System Identification

Like any methodological approach, control theory can only be as effective the information it is provided. If the system to which one is applying control theory fails to capture the important dynamics of the real-world psychological system, the methods we have presented here will be ill-equipped to inform intervention. Accordingly, how one decides to model a psychological process will have major consequences for control theory applications.

In the network psychometric literature, the most commonly used model, the VAR model, assumes stationarity: the means and variances of a time series are stable across time. In this paper, we used the VARI model, assuming that timeseries of symptom values are non-stationary, but that the change in symptom values from timepoint to timepoint is stationary. This ensured that the symptoms followed a unit-root process, allowing their trajectories to be permanently altered by interventions. This is vital, as a stationary VAR process is inherently mean reverting: no amount of intervention will permanently shift the expected symptom values away from the model-defined mean. Future research in the area of longitudinal psychometric network modeling should focus on modeling systems of symptoms in a way that allows disturbances (e.g., interventions) to impact the data-generating process directly. Such work will be especially important for researchers aiming to inform treatment. A state-space modeling approach is particularly amenable to this, as non-stationary processes can be modelled at the latent state level, while stationary error processes can be modeled at the indicator level. State-space modeling also allows for non-linear processes, an area where control theorists have developed several approaches for optimizing interventions.

An alternative approach to system identification than the one presented here would work not from data-based models, but from theory-based models. In recent years, several researchers have argued for the benefits of formalizing theories in clinical psychology as mathematical or computational models (e.g. Burger et al., 2019; Fried, 2020; Haslbeck, Ryan, Robinaugh, Waldorp, & Borsboom, 2019; Robinaugh, Haslbeck, et al., 2019). These formal theories often take the form of a causal network that aim to capture the dynamic processes that give rise to psychopathology. Although we herein demonstrated the use of control theory methods on networks derived from real data, there is no reason why the methods we used could not be applied to understanding the dynamics of a formal model. Indeed, we consider this to be a novel and promising application of control theory for psychological science. However, just as with models derived from empirical data, it is important to remember that the application of control theory will only be successful to the extent that the formal theory is an adequate representation of the target system of interest. Accordingly, it will be critical for theorists to develop strong formal theories if this approach is to be successful.

5.2.2. *The Importance of Measurement and Measurement Error*

The systems modeled in this paper were assumed, for the sake of simplicity, to have a white noise error component. White noise error processes do not impact the optimal control estimates for linear systems (Molenaar, 1987) and more complex error processes (e.g., the temporally correlated noise commonly found in psychological data) can be addressed through the use of stochastic control theory. However, there is another concern for psychological data is that it is important to consider and requires careful consideration to address: the nature of the error process. For symptoms that are theoretically stable across time, is the error a process that occurs at every timepoint and impacts the symptom value or does it instead simply reflect general sampling error? An error process that does impact symptom values (or another process of interest) is called a *disturbance*, and explicit modeling of disturbances is vital for understanding how the dynamical system underlying a given psychological disorder operates. Modeling disturbance is important not only in analyses, but also throughout data collection. For studies to accurately model symptom dynamics, they must also collect data on relevant events that could impact the symptoms under study (e.g., interpersonal conflicts). Only with that data collected can researchers disentangle true changes in the system due to disturbance vs. changes in the system due to sampling error.

Error also plays an important role in the domain of measurement. Studying how measurement error impacts statistical models of psychological data is a time-honored tradition in psychometrics. Control theory should be no exception. Some control theory literature does touch on measurement error (i.e. Schultz, 1964), but measurement error is particularly salient for clinical psychopathology research. Most recent applications of control theory to human behavior have observable outcomes (e.g., walking measured by actigraphy), and are less susceptible to measurement error. The measurement of psychological symptoms however typically relies on self-reported assessments of symptom severity/frequency, which are more susceptible to error, making it vital to evaluate effects of measurement error on control outcomes. One promising but currently underutilized way to account for measurement error in network psychometric is latent variable network modelling (Epskamp, Rhemtulla, & Borsboom, 2017), which should be further developed for longitudinal network modelling of non-stationary processes.

In addition to error in symptom measurement, there also likely to be error in intervention strength measurement. The strength of an intervention is not simply the immediate impact on symptoms, but also needs to take into account how the intervention interacts with subject specific dynamics. The consequences of this type of measurement error are especially salient in our present context, as psychological interventions have direct impacts on the lives and health of patients. The issue of how best to measure intervention strength is a relatively novel one for clinical psychology, but one that is critical to evaluating the expected impact of the intervention on systems of interest. Future control theory applications should examine how measurement error in the estimate of system dynamics, control inputs, and symptom observation during intervention may all bias conclusions.

5.2.3. The Future of Control Theory in Psychology

In this paper we aimed to provide an accessible introduction to control theory by focusing on a relatively straightforward set of analyses. However, the field of control theory is much broader than just these analyses and more complex analyses are available to help address challenges that are likely to be faced in applying this work within clinical psychology. For example, the analyses performed in this paper restricted symptoms values to a particular range and required interventions to be positive. Binary control inputs, or control inputs that have bounds, are referred to in the control theory literature as *saturated control inputs*. These restrictions led to interpretable trajectories and intervention regimes above, but can have large impacts on system behavior and are thus not completely optimal. When applying control theory to psychology, inputs and outputs will likely always be bounded/saturated, which necessitates more complex methods than we have used here when optimal performance is desired.

The constraint issue can be resolved using *hybrid model predictive control* (HMPC; Bemporad & Morari, 1999), which not only allows for constraints on both control inputs and state values, but also for categorical control inputs and state values. HMPC also brings with it the advantages of model predictive control, in that it can be applied to non-linear systems and allows for non-quadratic forms of the cost function. However, model predictive control approaches do not guarantee an *optimal* sequence of interventions and systems as it can be computationally difficult to determine the control laws for many variables (Kouvaritakis & Cannon, 2015). Nonetheless, psychological data is likely to be more amenable to model predictive control in practice, because there may be less need for rapid online control calculations such as in an engineering context. This would avoid some of the computational issues associated with applying this approach to high-dimensional data with a high sampling frequency.

Another challenge likely to be faced in clinical psychology is the possibility for time-delayed effects. Our interventions effects were modeled as instantaneous, with an intervention at time t causing a change in symptoms at time t . In reality, psychological interventions may not have such instantaneous effects. For example, an intervention reducing grief-related avoidance behavior may take weeks before it is accepted and adhered to by a patient. Such time-delayed effects, like saturated inputs, change how optimal controls should be calculated. Future research in this area should focus on bringing time-delayed control methods into control theory applications in psychology, as this would improve the external validity of the approach.

Perhaps the most important step in applying control theory within clinical psychology will be to move toward applying these methods within person-specific networks. In Section 4.1 we estimated the intervention effect at the group-level, assuming homogeneity in treatment response. Although interventions will likely have at least some common effect, the response of any single participant to any single intervention is likely heterogeneous. One solution is to explicitly estimate individualized treatment effects. However, there are several important considerations to note. First, the amount of data per person needed to estimate individual treatment effects is far greater than that needed when estimating group-level treatment. Second, estimation of an individual intervention effect requires collecting data before treatment for that given individual, which complicates optimally controlling the intervention applications for that person. Nonetheless, modeling individual treatment effects and applying them in a control setting could

potentially improve the personalization of psychological interventions.

Finally, we would like to stress the importance of cross-validation in future applications of control theory that compares simulated and empirical interventions. The methods and results presented in this paper take a simulation approach to evaluating the implied efficacy of interventions on psychological networks. This allows for formalized predictions of treatment response, without having to evaluate the treatment in an empirical study. A vital next step would be to compare these formalized predictions from the simulations to the actual treatment response in a new intervention study with a different sample of individuals. As mentioned above, this enables researchers to cross-validate their interventions. Any mismatch between the simulation-implied treatment effect and the empirical (observed) treatment effect would help reveal issues in modeling, data collection, or other threats to inference.

5.3. Conclusion

Control theory holds great promise for improving the development and evaluation of psychological interventions. Control theory as a field is more generally well developed; as such, future applications of control theory to psychology can benefit from many already existing tools. Although there are some unique challenges that psychological data also present, solving these challenges offers new opportunities to provide clinical science with more accurate predictions and stronger interventions that account for heterogeneity in both treatment and treatment response. In particular, control theory opens up new possibilities for personalized treatment strategies, where interventions could be applied based on the unique dynamics of a given patient. More work is clearly necessary, but we hope that applying control theory to the intervention sciences will ultimately improve people's real-life experiences and outcomes across a diversity of interventions, be they medical treatments, mental health management, or behavioral interventions.

References

- Åström, K. J. (2006). *Introduction to stochastic control theory*. Mineola, N.Y: Dover Publications.
- Bemporad, A., & Morari, M. (1999, March). Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35(3), 407–427. doi: 10/b765r5
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1), 289–300. doi: 10/gfpxdx
- Borsboom, D. (2008). Psychometric perspectives on diagnostic systems. *Journal of clinical psychology*, 64(9), 1089–1108.
- Borsboom, D. (2017). A network theory of mental disorders. *World Psychiatry*, 16(1), 5–13. doi: 10/f9ntpv
- Bringmann, L. F., Vissers, N., Wichers, M., Geschwind, N., Kuppens, P., Peeters, F., ... Tuerlinckx, F. (2013). A Network Approach to Psychopathology: New Insights into Clinical Longitudinal Data. *PLoS ONE*, 8(4), e60188. doi: 10/f4sxst
- Burger, J., van der Veen, D. C., Robinaugh, D., Quax, R., Riese, H., Schoevers, R. A., & Epskamp, S. (2019). *Bridging the Gap Between Complexity Science and Clinical Practice by Formalizing Idiographic Theories: A Computational Model of Functional Analysis* (Preprint). PsyArXiv. doi: 10.31234/osf.io/gw2uc
- Carver, C. S., & Scheier, M. F. (1998). *On the Self-Regulation of Behavior* (1st ed.). Cambridge University Press.
- Collins, L. M., Murphy, S. A., & Bierman, K. L. (2004, September). A conceptual framework for adaptive preventive interventions. *Prevention Science: The Official Journal of the Society for Prevention Research*, 5(3), 185–196. doi: 10/bpq4tq
- Contreras, A., Nieto, I., Valiente, C., Espinosa, R., & Vazquez, C. (2019). The Study of Psychopathology from the Network Analysis Perspective: A Systematic Review. *Psychotherapy and Psychosomatics*, 88(2), 71–83. doi: 10/gf8mtk
- Cramer, A. O. J., Waldorp, L. J., van der Maas, H. L. J., & Borsboom, D. (2010). Comorbidity: A network perspective. *Behavioral and Brain Sciences*, 33(2-3), 137–150. doi: 10/bwkc7q
- Dablander, F., & Hinne, M. (2019, May). Node centrality measures are a poor substitute for causal inference. *Scientific Reports*, 9(1), 6846. doi: 10/ggkjbp
- Epskamp, S. (2017). *Network Psychometrics* (Unpublished doctoral dissertation). University of Amsterdam.
- Epskamp, S., Rhemtulla, M., & Borsboom, D. (2017, December). Generalized Network Psychometrics: Combining Network and Latent Variable Models. *Psychometrika*, 82(4), 904–927. doi: 10/gcmfjj
- Epskamp, S., Waldorp, L. J., Möttus, R., & Borsboom, D. (2018). The Gaussian Graphical Model in Cross-Sectional and Time-Series Data. *Multivariate Behavioral Research*, 53(4), 453–480. doi: 10/gghm3c
- Fisher, A. J., Medaglia, J. D., & Jeronimus, B. F. (2018). Lack of group-to-individual generalizability is a threat to human subjects research. *Proceedings of the National Academy of Sciences*, 115(27), E6106–E6115. doi: 10/gdngc3
- Fisher, A. J., Reeves, J. W., Lawyer, G., Medaglia, J. D., & Rubel, J. A. (2017). Exploring the

- idiographic dynamics of mood and anxiety via network analysis. *Journal of Abnormal Psychology*, 126(8), 1044–1056. doi: 10/gcnj4g
- Fried, E. I. (2020). Lack of theory building and testing impedes progress in the factor and network literature.
- Fried, E. I., & Nesse, R. M. (2015). Depression is not a consistent syndrome: An investigation of unique symptom patterns in the STAR*D study. *Journal of Affective Disorders*, 172, 96–102. doi: 10/f6vdgz
- Fried, E. I., van Borkulo, C. D., Cramer, A. O. J., Boschloo, L., Schoevers, R. A., & Borsboom, D. (2017). Mental disorders as networks of problems: A review of recent insights. *Social Psychiatry and Psychiatric Epidemiology*, 52(1), 1–10. doi: 10/gf2zv7
- Frumkin, M., Piccirillo, M., Beck, E. D., Grossman, J., & Rodebaugh, T. (2019). *Feasibility and Utility of Idiographic Models in the Clinic: A Pilot Study* (Preprint). PsyArXiv. doi: 10.31234/osf.io/m34aw
- Hamdan, A. M. A., & Nayfeh, A. H. (1989). Measures of modal controllability and observability for first- and second-order linear systems. *Journal of Guidance, Control, and Dynamics*, 12(3), 421–428. doi: 10/dth34n
- Haslbeck, J. M. B., Ryan, O., Robinaugh, D. J., Waldorp, L., & Borsboom, D. (2019). *Modeling Psychopathology: From Data Models to Formal Theories* (Preprint). PsyArXiv. doi: 10.31234/osf.io/jgm7f
- Henry, T. R. (2020). *Netcontrol: Control theory methods for networks*. Retrieved from <https://CRAN.R-project.org/package=netcontrol> (R package version 0.1)
- Hyland, M. E. (1987). Control theory interpretation of psychological mechanisms of depression: Comparison and integration of several theories. *Psychological Bulletin*, 102(1), 109–121. doi: 10/fxcsw6
- Johnson, R. E., Chang, C.-H., & Lord, R. G. (2006). Moving from cognition to behavior: What the research says. *Psychological Bulletin*, 132(3), 381–415. doi: 10/cv92m3
- Jordan, D. G., Winer, E. S., & Salem, T. (2020, September). The current status of temporal network analysis for clinical science: Considerations as the paradigm shifts? *Journal of Clinical Psychology*, 76(9), 1591–1612. doi: 10/ggvwb5
- Kendler, K. S., Zachar, P., & Craver, C. (2011). What kinds of things are psychiatric disorders? *Psychological medicine*, 41(6), 1143–1150.
- Kim, N. S., & Ahn, W.-k. (2002). Clinical psychologists' theory-based representations of mental disorders predict their diagnostic reasoning and memory. *Journal of Experimental Psychology: General*, 131(4), 451–476. doi: 10/c678jm
- Kouvaritakis, B., & Cannon, M. (2015). *Model predictive control*. New York, NY: Springer Berlin Heidelberg.
- Lenferink, L. I., & Eisma, M. C. (2018, March). 37,650 ways to have “persistent complex bereavement disorder” yet only 48 ways to have “prolonged grief disorder”. *Psychiatry Research*, 261, 88–89. doi: 10/gc6wtj
- Levine, S. Z., & Leucht, S. (2016, December). Identifying a system of predominant negative symptoms: Network analysis of three randomized clinical trials. *Schizophrenia Research*, 178(1-3), 17–22. doi: 10/f9phk7
- Lewis, F. L., Vrabie, D. L., & Syrmos, V. L. (2012). *Optimal control* (3rd ed ed.). Hoboken:

Wiley.

- Ljung, L. (1999). *System identification: Theory for the user* (2nd ed ed.). Upper Saddle River, NJ: Prentice Hall PTR.
- Molenaar, P. C. (1987). Dynamic assessment and adaptive optimization of the pschotherapeutic process. *Behavioral Assessment*, 9(4), 389–416.
- Molenaar, P. C. (2010). Note on optimization of individual psychotherapeutic processes. *Journal of Mathematical Psychology*, 54(1), 208–213. doi: 10/cscfmz
- Nahum-Shani, I., Smith, S. N., Spring, B. J., Collins, L. M., Witkiewitz, K., Tewari, A., & Murphy, S. A. (2018, May). Just-in-Time Adaptive Interventions (JITAI)s in Mobile Health: Key Components and Design Principles for Ongoing Health Behavior Support. *Annals of Behavioral Medicine*, 52(6), 446–462. doi: 10/gdsk2p
- Olshevsky, A. (2014a, May). Minimal Controllability Problems. *arXiv:1304.3071 [cs, math]*.
- Olshevsky, A. (2014b, September). Minimum Input Selection for Structural Controllability. *arXiv:1407.2884 [cs, math]*.
- Pasqualetti, F., Zampieri, S., & Bullo, F. (2014). Controllability metrics, limitations and algorithms for complex networks. In (pp. 3287–3292). IEEE. doi: 10/ggkhs9
- Preumont, A. (1997). Controllability and observability. In *Vibration control of active structures: An introduction* (pp. 173–195). Dordrecht: Springer Netherlands. doi: 10.1007/978-94-011-5654-7%0089
- Radden, J. (2018). Rethinking disease in psychiatry: Disease models and the medical imaginary. *Journal of Evaluation in Clinical Practice*, 24(5), 1087–1092. doi: 10/ggkhs9
- Rivera, D. E., Hekler, E. B., Savage, J. S., & Downs, D. S. (2018). Intensively Adaptive Interventions Using Control Systems Engineering: Two Illustrative Examples. In L. M. Collins & K. C. Kugler (Eds.), *Optimization of Behavioral, Biobehavioral, and Biomedical Interventions: Advanced Topics* (pp. 121–173). Cham: Springer International Publishing. doi: 10.1007/978-3-319-91776-4_5
- Rivera, D. E., Pew, M. D., & Collins, L. M. (2007). Using engineering control principles to inform the design of adaptive interventions: A conceptual introduction. *Drug and Alcohol Dependence*, 88, S31–S40. doi: 10/cbtczz
- Robinaugh, D. J., Haslbeck, J. M. B., Waldorp, L., Kossakowski, J. J., Fried, E. I., Millner, A., ... Borsboom, D. (2019). *Advancing the Network Theory of Mental Disorders: A Computational Model of Panic Disorder* (Preprint). PsyArXiv. doi: 10.31234/osf.io/km37w
- Robinaugh, D. J., Hoekstra, R. H. A., Toner, E. R., & Borsboom, D. (2019). The network approach to psychopathology: A review of the literature 2008–2018 and an agenda for future research. *Psychological Medicine*, 1–14. doi: 10/gghndf
- Schultz, P. R. (1964). An Optimal Control Problem with State Vector Measurement Errors. In *Advances in Control Systems* (Vol. 1, pp. 197–243). Elsevier. doi: 10.1016/B978-1-4831-6717-6.50010-3
- Shear, M. K. (2010). Complicated grief treatment: The theory, practice and outcomes. *Bereavement Care*, 29(3), 10–14.
- Shumway, R. H., & Stoffer, D. S. (2017). *Time series analysis and its applications: With R examples* (4th ed.). New York, NY: Springer Science+Business Media.

- Sinclair, K., & Molenaar, P. C. (2008). Optimal control of psychological processes: A new computational paradigm. In *Bulletin de la Societe des Sciences Medicales Luxembourg* (Vol. 1, pp. 13–33). Luxembourg.
- Summers, T. H., Cortesi, F. L., & Lygeros, J. (2016). On Submodularity and Controllability in Complex Dynamical Networks. *IEEE Transactions on Control of Network Systems*, 3(1), 91–101. doi: 10/ggkhs5
- Wilcoxon, F. (1945). Individual Comparisons by Ranking Methods. *Biometrics Bulletin*, 1(6), 80. doi: 10/fnmmt8
- Young, G., Lareau, C., & Pierre, B. (2014). One Quintillion Ways to Have PTSD Comorbidity: Recommendations for the Disordered DSM-5. *Psychological Injury and Law*, 7(1), 61–74. doi: 10/gfddmn
- Zimmerman, M., Ellison, W., Young, D., Chelminski, I., & Dalrymple, K. (2015). How many different ways do patients meet the diagnostic criteria for major depressive disorder? *Comprehensive Psychiatry*, 56, 29–34. doi: 10/f7bcp5