

Structural Change and Homeostasis in Organizations: A Decision-Theoretic Approach^{*†}

Carter T. Butts

Department of Social and Decision Sciences
Center for the Computational Analysis of Social and Organizational Systems
Carnegie Mellon University

Kathleen M. Carley

Department of Social and Decision Sciences
Center for the Computational Analysis of Social and Organizational Systems
H.J. Heinz III School of Public Policy and Management
Carnegie Mellon University

v0.1, January 10, 2002

This is a draft document. Please do not copy or cite.

Abstract

We present here a decision-theoretic framework for the analysis of organizational change under risk. An algorithm is demonstrated which identifies optimal change paths given uncertainty involving execution time, intervention cost, and payoffs resulting from particular structural configurations. An elaboration of the basic framework to accommodate external structural perturbations is shown, and is applied to the problem of organizational homeostasis. Finally, an extension of the decision model is provided which admits multiple decision makers with divergent preferences and capacities for inducing organizational change.

Keywords: organizational design, structural change, network evolution, decision theory, homeostasis

^{*}This work was supported in part by the Office of Naval Research (ONR), United States Navy Grant No. 1681-1-1001944, by the Army Research Labs on the grant "Personnel Turnover and Team Performance," by the National Science Foundation under both Grant No. ITR/IM IIS-0081219 and the CASOS IGERT grant, and by the center for Computational Analysis of Social and Organizational Systems at Carnegie Mellon University. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Office of Naval Research, the National Science Foundation, or the U.S. government.

[†]The authors would like to thank David Rode for his helpful comments regarding this paper.

Organizations are not unitary, undifferentiated entities: each is composed of an assortment of lower level components (human actors, physical resources, tasks to be performed, etc.), components which are in some particular configuration with respect to one another at any given time. That organizational performance is contingent (at least in part) on the configuration of its parts has been understood at least since the work of Thompson (1967), and subsequent research (e.g., Perrow (1970); Pfeffer and Salancik (1978)) has continued to develop this theme. Given that organizational structures are relevant to concrete questions of performance and survival, the objective of treating structure as an explicit part of the organizational design process (Galbraith, 1977; Simon, 1973; Marschak, 1986) is an important one for organization theory.

While past research in this area has focused on identifying optimal organizational forms (Williamson, 1975; Duncan, 1979; Burton and Obel, 1984; Baligh et al., 1990; Crowston, 1994) or structural properties which correlate with performance generally (Blau, 1972; Mackenzie, 1978; Lincoln et al., 1986; Lin, 1994), our focus is not on where organizations should be going (so to speak), but rather on how they are to get there. In particular, we propose to consider the *process* of organizational change from an explicitly decision-theoretic point of view. Given a current form, a desired form, and a set of possible structural interventions, what is the optimal means of reaching the latter state? By formulating this question as a formal decision problem, we are able to obtain precise solutions under a reasonably wide range of initial assumptions.

The ultimate purpose of this exercise is two-fold. First, by demonstrating how the optimal change problem may be formulated (and, in some cases, solved) as a decision problem, we hope to contribute to the normative (i.e., social engineering) literature on the problem of organizational design. Given the substantial economic, physical, and social resources consumed by modern organizations, the difficult question of how they may be structured so as to meet particular needs is of obvious importance. When it is necessary for organizations to change form – as, indeed, occurs with considerable frequency – this process, too, begs formal examination.

Our second purpose is to contribute to the descriptive (i.e., social scientific) literature on change in organizations, by providing a baseline against which to compare actual change processes. Just as the rational actor model has proven consistently useful in the psychology of individual judgment and decision making (Tversky and Kahneman, 1986; Dawes, 1996), so too have optimization-based models contributed to the theory of organizations (Arrow, 1974; Williamson, 1975; Glance and Huberman, 1994). By examining deviations from optimal change behavior, we may illuminate the sorts of processes which govern real structural shifts; similarly, by means of our decision model, we may predict the effects of heuristic decision processes on organizational performance.

1 Concepts and Definitions

Before proceeding to the decision problem itself, we begin by setting out some basic concepts which will prove useful in the discussion to follow. These consist of representations for the organizational structures themselves, the transformations which will be used to take one structure into another (i.e., interventions), and payoffs to the Organizational Designer, respectively. Throughout, our approach will be to emphasize generality over problem-specificity; by maintaining a high level of abstraction, we hope to provide the widest possible scope for the methods described herein.

1.1 Organizational Structure

In order to provide a decision-theoretic framework for organizational change, we must first set forth a concrete representation of organizational structure. Following prior research in this area (Carley and Krackhardt, 1998; 1999), we take the structure of any given organization, ζ , as consisting of a collection of graphs; we make no further assumptions about the nature of the structures in question (e.g., whether the graphs are simple, directed, valued, etc.).¹ Such a structure is, in turn, taken to be a member of a much larger set of structures which could in principle be occupied by ζ . More formally:

Definition 1. Given an organization ζ composed of graphs G_1, \dots, G_n , we refer to $\mathbf{G} = \{G_1, \dots, G_n\}$ as the *structure set* of ζ . The set of all realizable \mathbf{G} is in turn represented by \mathbb{G} , the *structural universe* of ζ .

Note that by “realizable” structure sets, we mean those which are in fact available to the organization, subject to its definitional constraints. Thus, if we defined organizational structure purely in terms of reporting and task assignment relations, then all elements of \mathbb{G} would be constrained to have two members. For convenience, we may also (without loss of generality)² take all $G \in \mathbf{G}$ to have the same vertex set (though vertex sets may of course vary with \mathbf{G}). Though this last is not a requirement for the framework considered here, it does simplify matters considerably, and is recommended in practice.

1.2 Structural Transformations

Having established a set of potential structures from which the current state of an organization may be drawn, we now require a generic means of describing changes in that state. This is accomplished via the notion of the *structural transformation*, which takes the set of structures into itself:

¹In fact, even this assumption is not required for any of the formal results shown here; these depend only on the properties of the constructs \mathbb{G} , \mathbb{T} , and π (see below). Our discussion of potential structural payoff functions and our examples, however, do assume this.

²Vertices intrinsically uninvolved in a given relation (e.g., knowledge elements in a task assignment network) can simply be regarded as isolates for that relation.

Definition 2. Any $T : \mathbb{G} \rightarrow \mathbb{G}$ is said to be a *structural transformation* for the structural universe \mathbb{G} . The set of all realizable, distinct structural transformations is denoted by the *transformation universe* \mathbb{T} .

Once again, the “realizability” of the elements of \mathbb{T} is taken to refer to their empirical feasibility in the context of the structural universe at hand. Thus, in a typical case, realizable structural transformations might include the expansion or contraction of the vertex set (e.g., hiring/acquisition or firing/liquidation), edge addition or deletion (e.g., task or personnel reassignment), etc. For the moment, we will consider \mathbb{T} as unitary, in the sense that it represents the capability of a single abstract decision maker (the hypothetical “Organizational Designer”) to alter organizational structure. Later, we shall complicate this picture somewhat by assuming that transformations may arise from more than one source, one of the most important being Chance.

The above gives us a means of discussing the transition from one organizational form to another, but omits the important consideration of *time*. In general, structural change is not instantaneous: between the instant at which a change is decided upon and the moment at which it is realized, falls a period of implementation. This delay is significant for organizational decision makers, since time spent in a suboptimal configuration can result in lost earnings, higher production costs, or worse. Taking this into account, we introduce a function which takes as input a given transformation, returning the time required for said transformation to be implemented. Formally:

Definition 3. The time interval between the initiation of a structural transformation T and its completion (i.e., the realization of $T(\mathbf{G})$) is referred to as the *implementation time* of T , and is denoted $\Delta(T)$.

(Note that we assume here that implementation time is a function of the structural transformation, and not of the structure on which this transformation operates.)

Before moving ahead to a consideration of payoffs, we pause to combine the ideas of Definitions 1 and 2 into a single formalism. This formalism provides a useful shorthand for discussing the structural changes – particularly sequences of such changes – and will be used extensively in the discussion which follows. Taking the set of possible structures (the structural universe) together with the set of feasible transformations (the transformation universe), we form a multigraph which encodes the complete set of changes which could be made by the organization in question. By focusing on properties of this graph – rather than on the details of the particular elements of \mathbb{G} and \mathbb{T} – we can more easily study the properties of potential structural changes.

Definition 4. Given a transformation T and structural universe \mathbb{G} , the *change graph* (denoted G_T) of T on \mathbb{G} is the graph formed by $(\mathbb{G}, \{(\mathbf{G}, T(\mathbf{G})) \mid \mathbf{G} \in \mathbb{G}\})$. The collection of such graphs for all $T \in \mathbb{T}$ is then the *change multigraph* of \mathbb{T} on \mathbb{G} , denoted $\mathbf{G}_{\mathbb{T}}$. A walk from \mathbf{G}_α to \mathbf{G}_ω in $\mathbf{G}_{\mathbb{T}}$ is said to be a *change walk*, and if all vertices of such a walk are distinct, the walk is further said to be a *change path*.

To foreshadow what we shall see in the coming sections, it is the change walk (or, ultimately, the change path) which will be at the center of our attention. Intuitively, the change walk corresponds to a particular strategy for converting an organization from an initial form (\mathbf{G}_α) to a final one (\mathbf{G}_ω) by means of a series of concrete interventions. Likewise, each such walk also corresponds to a series of intermediate positions occupied by the organization during the change process: observe that any change walk from \mathbf{G}_α to \mathbf{G}_ω can be written either in terms of a vertex sequence (i.e., $\mathbf{G}_\alpha, \mathbf{G}_\beta, \dots, \mathbf{G}_\omega$) or in terms of an initial point together with a series of transformations (i.e., $\mathbf{G}_\alpha, T_1, \dots, T_n$).³ This follows from the fact that $\mathbf{G}_\alpha, \mathbf{G}_\beta, \dots, \mathbf{G}_\omega = \mathbf{G}_\alpha, T_1(\mathbf{G}_\alpha), \dots, T_n(\dots(T_1(\mathbf{G}_\alpha)))$ (itself a consequence of the definition of $\mathbf{G}_\mathbb{T}$). This simple duality will be of use for the arguments which follow.

1.3 Payoff Functions

In the context of an organizational design problem, it is assumed that there exists some decision maker who A) is able to act in altering organizational structure, and B) possesses well-formed preferences regarding the structures in question. While real-world organizations generally possess multiple stakeholders whose preferences are not entirely consistent with one another, we will initially limit ourselves to the simpler case. To represent this decision maker's preferences, we employ a *payoff function* whose value is to be maximized. This function may be either of transformations or of structural configurations, as indicated below:

Definition 5. Let $\pi(T)$ represent the payoff to the decision maker associated with employing transformation T , and let $\pi(\mathbf{G})$ be defined such that the total payoff to the decision maker associated with the organization ζ occupying state \mathbf{G} for duration t is equal to $t\pi(\mathbf{G})$ for all t, \mathbf{G} .

For the purposes of our analyses, we take all payoffs to be additive in transformations and time: thus the payoff associated with employing transformations T_1 and T_2 is $\pi(T_1, T_2) = \pi(T_1) + \pi(T_2)$, and the payoff resulting from occupying states \mathbf{G}_1 and \mathbf{G}_2 for durations t_1 and t_2 (respectively) would be $t_1\pi(\mathbf{G}_1) + t_2\pi(\mathbf{G}_2)$. Such an assumption seems a plausible first approximation, and buys considerable economy in terms of the tractability of the decision problem. Nonlinearly interacting payoff elements are left as a topic for future research.

It should be noted that an important task in defining the change decision problem is the identification of the payoff functions described above. Determination of $\pi(\mathbf{G})$, in particular, may be nontrivial. One potential candidate for such a payoff function may be derived as follows. Let f_1, \dots, f_m be a collection

³The astute reader will note that these representations are not *precisely* equivalent, in the sense that there may be more than one transformation connecting each pair of adjacent structures. That this discrepancy is of no consequence for the decision problem is shown below.

of structural indices such that $f_i : \mathbf{G} \rightarrow \mathbb{R} \forall \mathbf{G} \in \mathbb{G}$. Then, let f_1^*, \dots, f_m^* represent *optimal values* for these indices. Given this, a structural payoff function could reasonably take the form

$$\pi(\mathbf{G}) = \sum_{i=1}^m \beta_i g(f_i^* - f_i(\mathbf{G})) \quad (1)$$

where β is an *a priori* weight, and g is any decreasing even function (e.g., $-x^2$, $-|x|$). Such a payoff function corresponds to the idea that there exists an attribute space (provided by the indices) containing some point such that distance from said point is associated with decreasing payoffs. A related – and even simpler – family of functions is given by

$$\pi(\mathbf{G}) = \sum_{i=1}^m \beta_i g(f_i(\mathbf{G})) \quad (2)$$

which would be appropriate for linearly separable payoffs associated with particular indices. Personnel and resource costs, for instance, could easily be included in this manner (via vertex counting functions), as could certain types of transaction costs (e.g., via edge counts). Estimation of such functions from empirical performance data is fairly straightforward (provided such data are available).

Alternately, $\pi(\mathbf{G})$ can be determined by other means. Even in the absence of comprehensive performance data on existing organizations, sufficient knowledge of the organizational processes being assessed should permit estimation of $\pi(\mathbf{G})$ by means of simulation methods. (See, for instance, Lin (1994) for an archetypical example.) Given a candidate structure, \mathbf{G} , one can simulate said structure’s performance (possibly averaged over uncertain environmental conditions) in order to obtain an estimated payoff (see, e.g., Levitt et al. (1994)). Finally, it should be noted that the remaining aspects of the decision problem can still be treated in the absence of detailed specification of structural payoffs. Where the decision maker is indifferent to organizational structure, $\pi(\mathbf{G})$ can be set identically equal to zero; in this case, the optimal change problem simply becomes one of finding the least expensive transformations. Similarly, setting $\pi(\mathbf{G}) = -c$ (for some positive constant c) provides a very simple way to model costs due to implementation time.

Before moving on to the analysis of change decisions *per se*, a final note is in order regarding the payoffs associated with potential structural changes. Specifically, we require that no matter what the form of the payoff function, the final end state of the change process be preferred above any alternate structural position or series of transformations. While this may seem at first a non-obvious requirement, the need for it follows quite simply from the fact that our purpose is to identify the *optimal* walk from \mathbf{G}_α to \mathbf{G}_ω . If there exists some third structure, \mathbf{G}_β , such that $\mathbf{G}_\beta \succ \mathbf{G}_\omega$, then there is no such walk: we are better off staying at \mathbf{G}_β ! Similar logic applies in the case of transformation sequences.⁴

⁴E.g., it may become optimal to travel in endless loops.

We assume, by virtue of the fact that we have been asked to find the optimal \mathbf{G}_α , \mathbf{G}_ω walk, that our preferences are such that this is a sensible question. The specific requirement is given by the following axiom:

Axiom 1. *Given a decision to change organization ζ from structure \mathbf{G}_α to structure \mathbf{G}_ω , $\pi(\mathbf{G}_\omega) > \pi(\mathbf{G})$, $\Delta(T)\pi(\mathbf{G}_\omega) > \Delta(T)\pi(\mathbf{G}) + \pi(T) \forall T, \mathbf{G} : T \in \mathbb{T}, \mathbf{G} \in \mathbb{G}, \mathbf{G} \neq \mathbf{G}_\omega$.*

It should be emphasized that, in practice, Axiom 1 is quite easy to satisfy. If nothing else, simply putting a sufficiently large positive payoff on \mathbf{G}_ω will guarantee well-formed solutions; since the particular choice of optimal walk does not otherwise depend on this value, the exact number is irrelevant. Indeed, for most problems, it will suffice simply to declare $\pi(\mathbf{G}_\omega) = 0$, and to take all other payoffs in terms of losses relative to the destination state. So long as Axiom 1 is satisfied, the specific form of the payoff function may be chosen in accordance with the problem at hand.⁵

2 Riskless Change Decisions

Having set forth the basic elements of our decision-theoretic framework, we now proceed to put the pieces together in the context of a fairly simple problem: given a transformation universe, what is the optimal change walk between an initial structure set and some desired endpoint? In this simple case, we assume that implementation times and payoffs are constant, and rule out any exogenous perturbations to the organizational structure. As such, the decision is a riskless one, and merely requires us to search the space of walks for the element with the highest payoff.

We begin by computing total payoff associated with a change walk. Given structural and transformation universes \mathbb{G} and \mathbb{T} , let $\mathcal{W} = (\mathbf{G}_1, \dots, \mathbf{G}_n) = (\mathbf{G}_1, T_1, \dots, T_n)$ represent a change walk in $\mathbf{G}_\mathbb{T}$. From the above definitions, we may write the total payoff associated with \mathcal{W} as:

$$\pi(\mathcal{W}) = \Delta(T_1)\pi(\mathbf{G}_1) + \pi(T_1) + \dots + \Delta(T_n)\pi(\mathbf{G}_{n-1}) + \pi(T_n) \quad (3)$$

$$= \left(\sum_{i=1}^n \pi(T_i) \right) + \left(\sum_{i=1}^n \Delta(T_i) \pi(T_{i-1}(\dots(T_1(\mathbf{G}_1)))) \right) \quad (4)$$

Observe that Equation 3 can be partitioned into two components: the “direct” payoff associated with the application of structural transformations, and the indirect “exposure” payoff associated with the occupancy of intermediate states during the walk itself. Thus, it is not necessarily the case that a walk which is “cheap” in terms of low-cost transformations will be optimal overall. If a change walk causes the organization to spend long periods of time in a suboptimal state, the accumulated costs of exposure may overwhelm the savings from

⁵And, again, if the seemingly appropriate payoff function does not satisfy Axiom 1, this should suggest to the researcher that his or her problem may be improperly posed.

cutting corners on transformations; similarly, this same principle implies that the *order* in which changes are made may have a significant impact on the total payoff. For instance, an individual can be replaced by a dismissal followed by a new hire, or by a new hire followed by a dismissal of the original individual. The transformations involved are the same in both cases, but the optimal order will depend critically on the relative costs of duplicating versus being short on labor, and on the implementation time associated with hiring and dismissal processes. Where hiring is slow and the position in question is critical, it will often make sense to follow the latter procedure.

To summarize, then, the procedure for making an optimal riskless change decision is as follows:

1. Define \mathbb{G} , \mathbb{T} , Δ , π for the problem at hand (e.g., based on empirical data, first principles, etc.);
2. Identify the starting and ending points for the change walk, \mathbf{G}_α and \mathbf{G}_ω ;
3. Search the space of all change paths with the appropriate endpoints for a walk with the maximum payoff;
4. Any walk with the maximum payoff is an optimal change walk, and constitutes a solution to the change decision problem.

2.1 Properties of Optimal Change Walks

Before turning to the question of how optimal change walks may be identified, we pause momentarily to ask whether there are any more general statements we can make about the properties of such walks. Without knowing any further details of \mathbb{G} , \mathbb{T} , or π (beyond those already assumed), what can we say about $\mathbf{G}_\mathbb{T}$ and \mathcal{W} ? As it happens, even these constraints are sufficient to allow for some basic deductions regarding optimal change walks. These are interesting in and of themselves, in addition to being useful for proving the correctness of the algorithm which is to follow.

Our first result regards, in a sense, the “size” of the construct which is needed to solve the optimal change walk problem. Given the tremendous multiplexity of $\mathbf{G}_\mathbb{T}$ – that is, the potentially large number of edges between adjacent nodes – we are immediately led to wonder whether this profusion of ties is really necessary. Is there a way to “throw out” superfluous ties, so as to be left with a smaller (and in some respects, simpler) structure? As it happens, the answer is in the affirmative, as the following theorem demonstrates:

Theorem 1. *Given a payoff function π , $\forall \mathbf{G}_\mathbb{T} \exists$ a digraph $\mathbf{H}_\mathbb{T} : \mathbf{H}_\mathbb{T}$ contains an optimal change walk $\forall (\mathbf{G}_i, \mathbf{G}_j)$ having an optimal change walk in $\mathbf{G}_\mathbb{T}$, which is of equal total payoff.*

Proof. Given $\mathbf{G}_i, \mathbf{G}_j \in \mathbb{G}$, define $T^*(\mathbf{G}_i, \mathbf{G}_j) = T : T \in \mathbb{T}, T(\mathbf{G}_i) = \mathbf{G}_j, \pi(T) + \Delta(T)\pi(\mathbf{G}_i) = \max_{T': T'(\mathbf{G}_i) = \mathbf{G}_j} (\pi(T') + \Delta(T')\pi(\mathbf{G}_i))$. Then let $\mathbf{H}_\mathbb{T}$ be a

digraph whose vertex set is given by \mathbb{G} and whose edge set is given by $\{T^*(\mathbf{G}_i, \mathbf{G}_j) : \mathbf{G}_i, \mathbf{G}_j \in \mathbb{G}\}$.

We now show that all optimal change walks in $\mathbf{G}_{\mathbb{T}}$ are represented in $\mathbf{H}_{\mathbb{T}}$. Consider two vertices, $\mathbf{G}_\alpha, \mathbf{G}_\omega \in \mathbb{G}$, such that no $(\mathbf{G}_\alpha, \mathbf{G}_\omega)$ optimal change walk in $\mathbf{G}_{\mathbb{T}}$ belongs to $\mathbf{H}_{\mathbb{T}}$. For convenience, let us denote one of these walks by $\mathcal{W}^* = (\mathbf{G}_\alpha, T_1, \dots, T_n) = (\mathbf{G}_1, \dots, \mathbf{G}_n)$. Since $\mathcal{W}^* \not\subseteq \mathbf{H}_{\mathbb{T}}$, it follows that one of two conditions holds: i) $\exists T_i \in \mathcal{W}^* : \pi(T_i) + \Delta(T_i)\pi(\mathbf{G}_{i-1}) < \max_{T':T'(\mathbf{G}_i)=\mathbf{G}_j} (\pi(T') + \Delta(T')\pi(\mathbf{G}_i))$, or ii) $\pi(T_i) + \Delta(T_i)\pi(\mathbf{G}_{i-1}) = \max_{T':T'(\mathbf{G}_i)=\mathbf{G}_j} (\pi(T') + \Delta(T')\pi(\mathbf{G}_i)) \forall T_i \in \mathcal{W}^*$. Let us consider each in turn. If (i) is true, then it follows that $\pi(\mathcal{W}^*)$ can be increased by selecting an alternate transformation; since this contradicts the assumed optimality of \mathcal{W}^* , it follows that (i) is false. If (ii) is true, then (by construction) every adjacent vertex pair in is connected by an edge in $\mathbf{H}_{\mathbb{T}}$ which is of equal total value to the corresponding edge in \mathcal{W}^* . This implies, however, that the total value of the corresponding walk in $\mathbf{H}_{\mathbb{T}}$ is equal to $\pi(\mathcal{W}^*)$. Thus, this walk must be optimal, and since $\mathbf{H}_{\mathbb{T}} \subseteq \mathbf{G}_{\mathbb{T}}$, it must also belong to $\mathbf{G}_{\mathbb{T}}$. This contradicts our initial assumption, and thus it must be the case that all optimal change walks in $\mathbf{G}_{\mathbb{T}}$ are represented in $\mathbf{H}_{\mathbb{T}}$. \square

Intuitively, when moving between two adjacent vertices, we can effectively ignore any edges which are not of maximum local payoff; further, since any edges which remain must be of the same payoff, any one of these will be as effective as any other. Although we have focused on this result in its application to longer walks, the logic obviously holds generically for any optimal move. Expressed as a “rule,” the following corollary captures the key idea:

Corollary 1 (The “Local Domination” Rule). *In making a one-transformation move from \mathbf{G}_i to \mathbf{G}_j , it is always optimal to use a transformation with the highest immediate payoff.*

Proof. By construction, every edge of $\mathbf{H}_{\mathbb{T}}$ is of maximum immediate payoff (i.e., $\pi(T) + \Delta(T)\pi(\mathbf{G}) = \max_{T':T'(\mathbf{G}_i)=\mathbf{G}_j} (\pi(T') + \Delta(T')\pi(\mathbf{G}_i))$). By Theorem 1, $\mathbf{H}_{\mathbb{T}}$ also contains an optimal change walk for all pairs of vertices having a change walk in $\mathbf{G}_{\mathbb{T}}$, which is of equal total payoff to the optimal $\mathbf{G}_{\mathbb{T}}$ walks. Thus, it follows that it is never necessary to employ a transformation which does not have the highest immediate payoff. \square

Thus, as a practical matter, Corollary 1 contains the idea that an organization need never concern itself with local moves which are (locally) dominated. This is appealing, in that it suggests that a potentially large class of actions can be eliminated very quickly from consideration. As it turns out, there are other conceivable sets of actions which can be similarly ignored in attempting to make optimal changes. One such set is the set of transformation sequences which result in cycles, as is demonstrated by the following lemma:

Lemma 1. *All optimal change walks are change paths.*

Proof. Assume that there exists an optimal change walk from \mathbf{G}_α to \mathbf{G}_ω in $\mathbf{H}_\mathbb{T}$, denoted \mathcal{W}^* , which is not a change path. Then \mathcal{W}^* contains at least one cycle, \mathcal{C} , and at least one embedded $(\mathbf{G}_\alpha, \mathbf{G}_\omega)$ path, \mathcal{P} ; without loss of generality, choose \mathcal{P} and \mathcal{C} to be of maximum payoff. Since \mathcal{W}^* is optimal, it follows that

$$\pi(\mathcal{P} \cup \mathcal{C}) \geq \pi(\mathcal{P}) + \Delta(\mathcal{C})\pi(\mathbf{G}_\omega),$$

and hence that

$$\pi(\mathcal{C}) \geq \Delta(\mathcal{C})\pi(\mathbf{G}_\omega),$$

which, finally, implies

$$\sum_{i=1}^{|\{T:T \in \mathcal{C}\}|} (\pi(T_i) + \Delta(T_i)\pi(\mathbf{G}_{i-1})) \geq \sum_{i=1}^{|\{T:T \in \mathcal{C}\}|} \Delta(T_i)\pi(\mathbf{G}_\omega).$$

This is only possible if $\exists \mathbf{G} \neq \mathbf{G}_\omega, T : \pi(T) + \Delta(T)\pi(\mathbf{G}) \geq \Delta(T)\pi(\mathbf{G}_\omega)$, which contradicts Axiom 1. It therefore follows that all optimal change walks are change paths. \square

As before, the intuition is fairly straightforward. Given that (by Axiom 1) we would rather be at our destination than anywhere else, actions that artificially prolong our journey cannot be optimal. Clearly, walks which are not paths have this property, since they necessarily contain shorter sub-walks which still reach the desired goal. This observation, like that of Theorem 1, lends itself to a simple behavioral rule, namely:

Corollary 2 (The “No-Backsies” Rule). *It is never optimal for an organization which is changing from form \mathbf{G}_α to \mathbf{G}_ω to revert to a previous intermediate form.*

Proof. This follows immediately from Lemma 1, as any vertex sequence which contains repeated elements cannot (by definition) constitute a path, and since optimal change walks are always paths. \square

Simple as it is, the “no-backsies” rule does have some interesting implications. For instance, it suggests that firms which are observed to “churn” personnel via repeated cycles of hiring and firing are most likely not following an optimal change path. (Although this behavior may also be the result of external perturbations; see below regarding this point.) In general, intendedly cyclical behavior is diagnostic of suboptimality, and looking for the former may prove a useful heuristic for locating examples of the latter in the field.

2.2 Finding an Optimal Change Path

As we established with Theorem 1 and Lemma 1, the problem of finding an optimal change walk can be reduced to that of finding a maximum payoff path

on $\mathbf{H}_{\mathbb{T}}$. As it happens, this problem is itself simply an alternate form of the shortest-path problem, a well-known problem of combinatorial optimization. To see that this is so, one must merely reframe the payoff associated with the change walk in terms of the *cost* of the path relative to spending the entire traversal time at the destination state. (Recall that, by Axiom 1, the latter payoff must be strictly greater than that of any walk on $\mathbf{G}_{\mathbb{T}}$.) Then it follows that the path of minimum cost is that of maximum payoff, and an algorithm which solves the former problem can also be used to solve the latter.

Although, in principle, any shortest-path algorithm could be modified to solve the optimal change path problem (see Ahuja et al. (1993) and Cook et al. (1998) for in-depth discussions of shortest-path algorithms) we present here a variant of Dijkstra’s well-known label-setting algorithm (Dijkstra, 1959).⁶ This is shown in Algorithm 1. Although fairly typical in form, Algorithm 1 does contain some features which relate specifically to the optimal change problem. First, and most trivially, we exploit the fact that we are interested only in the $(\mathbf{G}_{\alpha}, \mathbf{G}_{\omega})$ path by terminating execution as soon as the shortest such path is found (see line 19). Second, as already noted, we have reversed the usual sense of the optimization by seeking a maximum payoff path instead of a path of minimum cost; this is operationalized by line 13, and is rendered feasible by Axiom 1. Third, we deal with the multiplexity of $\mathbf{G}_{\mathbb{T}}$ via the loop in lines 9-18, which ensures that only the maximum payoff (minimum cost) arc is used. Although this is presented here as a runtime calculation in order to conserve memory, it is also in principle possible to discard all but the least expensive edges in an initial step, thereby operating directly on $\mathbf{H}_{\mathbb{T}}$. This seems unlikely to be practical in most situations, but some performance gains may still be possible (depending on π and $|\mathbf{G}_{\mathbb{T}}|$) if some transformations could be eliminated *ex ante*. This leads us to our fourth significant modification to the basic algorithm: because it will rarely if ever be feasible to store distances to all members of \mathbb{G} , we attempt to minimize memory usage by storing distances only for those vertices we have visited. (Note that the default distance set by line 1 can be implemented via implicit storage (i.e., as a default value), with explicit storage only of visited vertices.) Indeed, storage is dominated by the distance and predecessor lists, and will be on the order of the number of vertices visited prior to finding the shortest path. Algorithm 1 is constructed so as to minimize the number of vertices visited, and hence should require as little storage as possible.

When executed, Algorithm 1 will yield the payoff of the optimal $(\mathbf{G}_{\alpha}, \mathbf{G}_{\omega})$ path as $d(\mathbf{G}_{\omega})$. The specific path followed may be reconstructed via the predecessor list, *pred*, which indicates both the predecessor and transformation employed for every visited vertex other than \mathbf{G}_{α} . (One simply starts with *pred*(\mathbf{G}_{ω}) and works backwards.) If *pred* is stored as a tree rather than a list, this operation can be performed linearly in the length of the optimal path. As a minor note, we have tacitly assumed throughout this discussion that a $(\mathbf{G}_{\alpha}, \mathbf{G}_{\omega})$ path actually *exists*; were this not so, the optimization question would be mean-

⁶The performance of Dijkstra’s algorithm is close to optimal for dense graphs (Ahuja et al., 1993), and thus it serves as an obvious starting point for the present application.

Algorithm 1 Identification of Optimal Change Paths

```
1:  $d(\mathbf{G}) := -\infty \forall \mathbf{G} \in \mathbf{G}_{\mathbb{T}}$ 
2:  $d(\mathbf{G}_{\alpha}) := 0$ 
3:  $vis := \{\mathbf{G}_{\alpha}\}$ 
4:  $pred(\mathbf{G}_{\alpha}) := \emptyset$ 
5:  $flag := \text{FALSE}$ 
6: while  $\neg flag, |vis| > 0$  do
7:    $\mathbf{G} := \mathbf{G}_i \in vis : d(\mathbf{G}_i) = \max_{\mathbf{G}_j \in vis} d(\mathbf{G}_j)$ 
8:    $vis := vis \setminus \mathbf{G}$ 
9:   for all  $T \in \mathbb{T}$  do
10:    if  $d(T(\mathbf{G})) = -\infty$  then
11:       $vis := vis \cup T(\mathbf{G})$ 
12:    end if
13:     $\delta := d(\mathbf{G}) + \pi(T) + \Delta(T) (\pi(\mathbf{G}) - \pi(\mathbf{G}_{\omega}))$ 
14:    if  $\delta > d(T(\mathbf{G}))$  then
15:       $d(T(\mathbf{G})) := \delta$ 
16:       $pred(\mathbf{G}_{\alpha}) := (\mathbf{G}, T)$ 
17:    end if
18:  end for
19:  if  $d(\mathbf{G}_{\omega}) > \max_{\mathbf{G}_i \in vis} d(\mathbf{G}_i)$  then
20:     $flag := \text{TRUE}$ 
21:  end if
22: end while
```

ingless! Still, in case the possibility of attaining an optimal solution should be in question, a check for the size of the “to visit” stack (*vis*) is performed at line 6. Should *vis* become empty while $d(\mathbf{G}_\omega) = -\infty$, this indicates that no $(\mathbf{G}_\alpha, \mathbf{G}_\omega)$ exists, and that the optimization problem is ill-posed.

We conclude our discussion of the optimal change algorithm with a proof of its correctness:

Theorem 2. *Algorithm 1 terminates with a solution to the optimal change problem, if such a solution exists.*

Proof. We begin our proof of the correctness of the algorithm by demonstrating that, at each iteration, \mathbf{G} is chosen so as to be “closest” (i.e., connected by a path of maximum payoff) to \mathbf{G}_α . (By Lemma 1, we may ignore any non-path walks.) Consider the first iteration, in which $\mathbf{G} = \mathbf{G}_\alpha$: by definition, \mathbf{G}_α is closest to itself. Now, consider the n th iteration, under the assumption that all vertices chosen so far have met the closeness condition. By step 7, the \mathbf{G} which is chosen is closest to \mathbf{G}_α . Further, by the finite loop of lines 9-18, all vertices which are adjacent to \mathbf{G} are placed at distance $d(\mathbf{G})$ plus their distance from \mathbf{G} . Plainly, this is an upper bound on their distances (a lower bound on the path of maximum payoff). If it is not also a lower bound for a given vertex, then it follows from the criterion of line 7 that there exists some other vertex which will be visited first, and which will lower the distance (raise the payoff) for said vertex (by line 14). Because every vertex adjacent to \mathbf{G} is added to *vis*, every such vertex is visited; from the forgoing, it follows that when this vertex is visited, it will be of maximum closeness/payoff. Therefore, it follows that the condition will be satisfied for these vertices as well, and, by induction, that the condition is satisfied for all vertices.

Given that every vertex adjacent to a visited vertex is visited (unless the termination condition has been met), the above implies that every vertex connected to \mathbf{G}_α is visited in descending order of payoff. To complete the proof, we note that the terminating conditions of lines 19 and 6 can only be met if every other path is of strictly lower payoff than an already uncovered $(\mathbf{G}_\alpha, \mathbf{G}_\omega)$ path, and/or if no vertices remain to be visited. If the latter is true, then there can exist no higher-payoff path than that which has been uncovered (if any): if and only if $\pi(\mathbf{G}_\omega) > -\infty$, such a path must exist. (Otherwise, there is no solution to find.) If the latter is false, then the observed $(\mathbf{G}_\alpha, \mathbf{G}_\omega)$ path is of higher value than any other path (by the induction argument above), and we may terminate the algorithm. The optimal change path may be derived by following the predecessors of \mathbf{G}_ω backwards until \mathbf{G}_α is reached. \square

3 Risky Change Decisions

In our discussion so far, we have mercifully avoided the vagueness and uncertainty associated with real-world decisions: we have presumed that everything from implementation times to the payoffs associated with particular structural

forms are fixed and known. Here, we generalize our earlier results to decisions in which implementation times and payoffs are uncertain.

Under an uncertainty assumption, implementation time and payoffs (especially the structural payoff) are taken to be random variables rather than constants. Thus, one cannot be sure precisely how long it will take to hire someone (for instance), or exactly what the effects will be of having lost the manager of a particular division. Introducing such uncertainty provides for one obvious source of risk in organizational design, and allows for decisions which are robust to unexpected delays.

For Δ, π random, the obvious optimization criterion is the expected payoff; applying the expectation operator to Equation 3 gives us

$$E(\pi(\mathcal{W})) = E(\Delta(T_1)\pi(\mathbf{G}_1)) + E(\pi(T_1)) + \dots \quad (5)$$

$$+ E(\Delta(T_n)\pi(\mathbf{G}_{n-1})) + E(\pi(T_n))$$

$$= \left(\sum_{i=1}^n E(\pi(T_i)) \right) + \left(\sum_{i=1}^n E(\Delta(T_i)\pi(T_{i-1}(\dots(T_1(\mathbf{G}_1)))))) \right) \quad (6)$$

Note that this result does *not* assume independence. If we *are* willing to take durations and payoffs as independent, we may obtain variances in the same manner:

$$Var(\pi(\mathcal{W})) = Var(\Delta(T_1)\pi(\mathbf{G}_1)) + Var(\pi(T_1)) + \dots \quad (7)$$

$$+ Var(\Delta(T_n)\pi(\mathbf{G}_{n-1})) + Var(\pi(T_n))$$

$$= \left(\sum_{i=1}^n Var(\pi(T_i)) \right) + \left(\sum_{i=1}^n Var(\Delta(T_i)\pi(T_{i-1}(\dots(T_1(\mathbf{G}_1)))))) \right) \quad (8)$$

These may be used to assess the level of risk associated with each change path. Although we assume risk-neutrality for the moment, the generalization of the approach to incorporate non-risk neutral preferences is an important direction for future development.

3.1 Finding an Optimal Change Path Under Risk

Finding an optimal change path under risk is not qualitatively distinct from the riskless case. By replacing transformation payoffs by their expectations, we may use a slightly modified version of Algorithm 1 (using Equation 5) to identify a change path which maximizes the expected path payoff. As before, simulation steps may need to be added to estimate expectations for payoffs associated with structural positions, the implementation of which will obviously be problem-specific. In general, however, the change from riskless to risky change decisions is fairly transparent in terms of the procedures involved.

4 Structural Perturbations

The foregoing has treated the problem of optimal change decisions in an environment in which the only source of structural change is the application of interventions by the Organizational Designer. While we may have been uncertain regarding actual payoffs and implementation times, we were nevertheless able to rely on the fact that the organizational structure itself was never in question. Although reasonable as a first approximation, we are naturally led to ask how things might differ if this constraint were relaxed. What if, in addition to our own interventions, organizational structure were subject to perturbations arising from other sources?

The core of our perturbation framework is as follows. Assume that, in addition to the Organizational Designer, Chance may attempt to modify organizational structure by applying transformations from the set \mathbb{T}_C ; \mathbb{T}_C need not be equal to \mathbb{T} . For the purposes of our analyses, we shall further presume that each player may elect to apply only one transformation at a time (i.e., before a response from the other player is possible), and that any transformation applied by the other player begins implementation immediately following the completion of the transformation which is already in progress (assuming one has already been applied). Thus, if the Designer applies transformation T_1 at time t , then Chance may respond with any transformation $T_C \in \mathbb{T}_C$ from time t to $t + \Delta(T_1)$. If Chance has not applied a transformation by time $t + \Delta(T_1)$, then the Designer may do so; otherwise, the Designer must wait until $t + \Delta(T_1) + \Delta(T_C)$ for his or her opportunity. If no transformation is “in play,” neither player is obligated to act, and either may wait for any length of time before attempting to apply a new transformation.

With respect to payoffs, we follow game-theoretic convention in presuming that Chance does not behave strategically; although its behaviors could conceivably depend in some way on those of the Organizational Designer,⁷ we obviously do not endow it with putative reasoning capability! Payoffs for the Designer are assumed to follow the same form as previously specified, without regard to whether the transformations employed were played by the Designer or by Chance. Since it is not required that $\mathbb{T}_C = \mathbb{T}$, this is without loss of generality.

4.1 Homeostasis: the Cost of Staying in Place

Homeostasis – in the sense of maintaining a fairly constant structural form – is an important element of organizational survival. An organization which is unable to maintain a state of homeostasis for long periods of time runs the risk of incurring sufficient costs to induce mortality (Hannan and Freeman, 1977), and simulation studies of organizations which are in constant flux suggest that survivors may still suffer from serious degradation in performance (Carley and Svoboda, 1996).

⁷For instance, applying a transformation which removes personnel may increase the probability that Chance will apply a transformation resulting in further personnel loss, thus providing a model for the “snowballing” of turnover processes (Krackhardt, 1986).

In any event, presuming a relatively stable environment, failure of homeostasis will almost always entail⁸ that an organization spend relatively large periods of time in sub-optimal structural states, and we may thus realistically expect it to impact performance.

In the context of our present study, an obvious question to ask regarding homeostasis is the cost associated with maintaining a particular form in the presence of external perturbations. This, of course, presumes that circumstances allow homeostasis to be maintained at all; depending on the extent of environmental disruption, this may or may not be the case. We thus begin our discussion of organizational homeostasis with the latter issue.

4.1.1 The Possibility of Homeostasis

Before moving to the question of homeostasis costs, we begin with an even more basic question: under what conditions can an Organizational Designer hope to maintain homeostasis at all? Consider an organization which is in structural equilibrium at state \mathbf{G}^* , and which is faced by structural perturbations which arrive as independent Poisson events. In particular, let us assume that Chance attempts to apply transformations via a Poisson process with constant intensity λ , with the particular transformation $T \in \mathbb{T}_C$ being selected by an independent multinomial draw with probability p_T ; if Chance already has a transformation “pending” or in the process of implementation, assume that the attempt is abandoned. Assuming that \mathbf{G}^* is the preferred organizational state, the optimal response from the Designer is to respond to any perturbation by attempting to restore equilibrium, i.e., to return to \mathbf{G}^* by the optimal change path. (We will consider here only cases in which such a return is possible; we assume that there exists no $T \in \mathbb{T}_C, \mathbf{G} \in \mathbb{G}$ such that a $(T_C(\mathbf{G}), \mathbf{G}^*)$ path does not exist in \mathbb{G}_T .) Given such a baseline scenario, what can be said about the possibility of maintaining homeostasis?

To get a sense of this, we begin by considering an upper bound on the perturbation rate (λ) compatible with homeostasis. Consider the set of transformations available to the Designer, \mathbb{T} . Clearly, no matter what perturbation Chance applies⁹, the Designer cannot restore equilibrium faster than $\Delta_{\min} = \min_{T \in \mathbb{T}} \Delta(T)$. Thus, if λ is sufficiently large for the inter-perturbation time to be small relative to Δ_{\min} , homeostasis becomes impossible. Given that perturbations occur as a Poisson process, this implies the condition

$$\frac{1}{\lambda} \gg \Delta_{\min} \tag{9}$$

⁸The exception being when there are sufficient numbers of optimal forms that the organization may constantly shift between them; this seems unlikely in practice, however.

⁹Other than the identity transformation; since including this within \mathbb{T}_C would be equivalent to a reduction in λ , we can safely ignore this possibility.

or, alternately,

$$\lambda \ll \frac{1}{\Delta_{\min}}. \quad (10)$$

To refine this constraint somewhat, we can also consider the expected time spent in homeostasis; that is, the expected time spent in state \mathbf{G}^* between perturbations. Following the above argument, a firm lower bound on the total re-equilibration time of the organization following a perturbation is given by Δ_{\min} . For convenience, let us refer to the inter-perturbation time as t_C , with $t_C \sim \exp(\lambda)$. Then a lower bound on the expected time spent in homeostasis must be given by

$$\varepsilon = 0p(t_C \leq \Delta_{\min}) + \int_{\Delta_{\min}}^{\infty} (t_C - \Delta_{\min}) \lambda e^{-\lambda t_C} dt_C \quad (11)$$

$$= -\left(t_C + \frac{1}{\lambda}\right) e^{-\lambda t_C} \Big|_{\Delta_{\min}}^{\infty} + \Delta_{\min} e^{-\lambda t_C} \Big|_{\Delta_{\min}}^{\infty} \quad (12)$$

$$= \left(0 + \left(\Delta_{\min} + \frac{1}{\lambda}\right) e^{-\lambda \Delta_{\min}}\right) + (0 - \Delta_{\min} e^{-\lambda \Delta_{\min}}) \quad (13)$$

$$= \frac{1}{\lambda} e^{-\lambda \Delta_{\min}}. \quad (14)$$

For a given bound, we can then identify the associated maximum λ by

$$\lambda \leq \frac{1}{\Delta_{\min}} W\left(\frac{\Delta_{\min}}{\varepsilon}\right), \quad (15)$$

where W is the Lambert-W function.¹⁰ One obvious value of interest is the λ bound associated with $\varepsilon = \Delta_{\min}$, i.e., the upper bound for the point at which the expected time in homeostasis is equal to the expected time spent out of equilibrium. In that case we have

$$\lambda \leq \frac{1}{\Delta_{\min}} W\left(\frac{\Delta_{\min}}{\varepsilon}\right) \quad (16)$$

$$= \frac{1}{\Delta_{\min}} W\left(\frac{\Delta_{\min}}{\Delta_{\min}}\right) \quad (17)$$

$$= \frac{W(1)}{\Delta_{\min}} \quad (18)$$

$$\approx \frac{0.567143209}{\Delta_{\min}}. \quad (19)$$

Thus, a somewhat more refined bound is given by approximately 57% of the reciprocal of the minimum re-equilibration time. For perturbation rates in

¹⁰The Lambert-W function is defined as the principal solution to the inverse of $f(W) = We^W$.

excess of this, *no* organization can reliably maintain homeostasis – regardless of the other properties of the structures or transformations involved.

The above provides a fairly general result, which gives us some intuition about the maximum perturbation rate at which an arbitrary organization can potentially maintain its form. Of course, the maximum rate for any *particular* organization may be quite a bit lower than this, depending on the properties of \mathbb{T} , \mathbb{T}_C , the p_T s, and \mathbf{G}^* . While no simple expression exists for this quantity, we can nevertheless obtain numerical estimates using the procedures already developed for optimal change paths. Let $\mathcal{P}^*(\mathbf{G}_i, \mathbf{G}_j)$ be the optimal change path from \mathbf{G}_i to \mathbf{G}_j in $\mathbf{G}_{\mathbb{T}}$, and let $\Delta(\mathcal{P}^*(\mathbf{G}_i, \mathbf{G}_j)) = \sum_{T \in \mathcal{P}^*(\mathbf{G}_i, \mathbf{G}_j)} \Delta(T)$ be the total duration of \mathcal{P}^* . Then an approximation to the expected time spent in homeostasis for an organization at \mathbf{G}^* is given by

$$\begin{aligned} \varepsilon \approx & \sum_{T \in \mathbb{T}_C} p_T [0p(t_C \leq \Delta(\mathcal{P}^*(T(\mathbf{G}^*), \mathbf{G}^*))) \\ & + \int_{\Delta(\mathcal{P}^*(T(\mathbf{G}^*), \mathbf{G}^*))}^{\infty} (t_C - \Delta(\mathcal{P}^*(T(\mathbf{G}^*), \mathbf{G}^*))) \lambda e^{-\lambda t_C} dt] \end{aligned} \quad (20)$$

$$\begin{aligned} = & \sum_{T \in \mathbb{T}_C} p_T \left[- \left(t_C + \frac{1}{\lambda} \right) e^{-\lambda t_C} \Big|_{\Delta(\mathcal{P}^*(T(\mathbf{G}^*), \mathbf{G}^*))}^{\infty} \right. \\ & \left. + \Delta(\mathcal{P}^*(T(\mathbf{G}^*), \mathbf{G}^*)) e^{-\lambda t_C} \Big|_{\Delta(\mathcal{P}^*(T(\mathbf{G}^*), \mathbf{G}^*))}^{\infty} \right] \end{aligned} \quad (21)$$

$$= \sum_{T \in \mathbb{T}_C} p_T \left[\frac{1}{\lambda} e^{-\lambda \Delta(\mathcal{P}^*(T(\mathbf{G}^*), \mathbf{G}^*))} \right]. \quad (22)$$

This value is an approximation of the true ε in that it does not consider any additional disruptive effect of perturbations which occur during the re-equilibration process. This effect will be small whenever $\sum_{T \in \mathbb{T}_C} p_T \Delta(\mathcal{P}^*(T(\mathbf{G}^*), \mathbf{G}^*)) \ll \frac{1}{\lambda}$, or (more generally) when the duration of the average perturbed optimal path is approximately equal to $\sum_{T \in \mathbb{T}_C} p_T \Delta(\mathcal{P}^*(T(\mathbf{G}^*), \mathbf{G}^*))$. In any event, since additional perturbations should, on average, increase the re-equilibration time, Equation 22 serves as an upper bound on ε .

Numerical estimation is necessary to obtain the λ corresponding to $\varepsilon = \sum_{T \in \mathbb{T}_C} p_T \Delta(\mathcal{P}^*(T(\mathbf{G}^*), \mathbf{G}^*))$, since Equation 22 cannot be solved directly. For small $|\mathbb{T}_C|$, this can reasonably be accomplished by using Algorithm 1 to find the duration of each perturbed path, and then solving for λ using an iterative search procedure (since the Δ s are independent of λ). Standard methods for such estimation can be found in any numerical methods text, e.g. Press et al. (1986).

4.1.2 Homeostasis Costs

As we have shown, the task of maintaining homeostasis is nontrivial; indeed, if perturbations occur with sufficient frequency, it may be difficult or impossible. Where perturbations are less frequent, the *ability* of the Designer to maintain equilibrium is not in question. The *cost* of maintaining equilibrium, however, is another matter entirely. If a high-performing structural configuration is surrounded (in the change multigraph formed by \mathbb{T}_C on \mathbb{G}) largely by low-performing configurations, the long-run average payoff associated with attempting to maintain this position may be severely reduced.

Here, we provide an expression for the approximate expected payoff associated with attempting to maintain an equilibrium position, \mathbf{G}^* . For convenience, we take all payoffs to be given relative to $\pi(\mathbf{G}^*)$; thus, the expression can be directly interpreted as the cost associated with homeostasis. As before, we assume Chance to apply transformations as Poisson events, with the transformation type selected via a multinomial process. Applying these assumptions then yields, for the expected homeostasis cost per unit time,

$$E(\pi^*(\mathbf{G}^*)) \approx \lambda \sum_{T \in \mathbb{T}_C} p_T (\pi(T) + \pi(\mathcal{P}^*(T(\mathbf{G}^*), \mathbf{G}^*))), \quad (23)$$

which converges almost surely to $E(\pi^*(\mathbf{G}^*))$ in the small- λ limit. (In practice, the approximation will be valid so long as $\sum_{T \in \mathbb{T}_C} p_T \Delta(\mathcal{P}^*(T(\mathbf{G}^*), \mathbf{G}^*)) \ll \frac{1}{\lambda}$.) For small $|\mathbb{T}_C|$, Equation 23 can be calculated very straightforwardly using Algorithm 1. If $|\mathbb{T}_C|$ is large, it may be necessary to resort to a more restricted sampling strategy; where the distribution of transformation probabilities (p_T) is extremely concentrated, one such option is to sample from \mathbb{T}_C using the transformation probabilities. This will tend to omit low-probability transformations, while still resulting in an unbiased estimate.

4.2 Optimal Changes Under Perturbation Risk

We have already considered the question of optimal change decisions in the context of perfect information, and of uncertainty regarding implementation time and payoffs. A far deeper challenge confronts us in facing the problem of finding an optimal change path under the risk of structural perturbations. Unlike the case of homeostasis, where our attention was focused primarily on small deviations from a steady state, we have now to consider the effect of such perturbations at any point along the change path. A complete solution to the optimal change problem under perturbation risk thus requires a *change strategy*, i.e., an algorithm which provides directions on how to proceed from each possible state towards the final objective. Unfortunately, such an algorithm is likely to be infeasible for any but the very smallest organizations. Thus, instead, we provide here some comments regarding a heuristic approach to this problem.

First, we note that while a complete solution is particularly difficult, it is nonetheless possible to obtain results regarding the robustness of various heuristics via simulation methods. The most obvious such heuristic is simply the

direct application of the standard optimal change path result; should the execution of this path be interrupted by a perturbation, then a new optimal path is computed based on the perturbed location, and this path is followed until the destination is reached or another perturbation ensues. To estimate the total payoff of such a strategy under structural perturbations, one would simply run a series of simulations in which the algorithm was applied, together with the perturbation process, to the problem. It is then a standard result (see e.g., Kalos and Whitlock (1986)) that the distribution of payoffs obtained will then converge to the true distribution. While this obviously does not improve upon the standard approach, it can nevertheless offer a test of the robustness of the basic strategy to external interference of various kinds.

Another variation on the above theme is the evaluation of alternative heuristics such as hill climbing, simulated annealing, and the like. Annealing processes, in particular, have been applied to the descriptive modeling of structural change in the past (Carley and Svoboda, 1996), and it may be useful to compare their performance to that of simple optimization. A more exotic alternative may be the use of genetic programs (Koza, 1992) to search the space of potential strategies for high-performing heuristics. Genetic programs have proven reasonably successful in finding solutions within large search spaces with complex correlation structures; thus, there is some reason to believe that they may perform here as well.

5 Structural Changes with Multiple Decision Makers

Thus far (with the rather mindless exception of Chance), our analyses have been confined to the decision problem faced by a single Organizational Designer who wishes to identify an optimal means of moving ζ from \mathbf{G}_α to \mathbf{G}_ω . One important extension of these ideas is to contexts in which multiple decision makers are involved in the change decision; these “structural games” include strategic issues which are absent in the “lone Designer” scenario. While there are many ways in which multiple decision makers could enter into the change process, we will here focus on two: a basic principal-agent game, and a simple bargaining problem. Although each is treated with some brevity, our aim is to highlight a few basic results which may be of reasonably broad applicability. Further investigation into the properties of structural games is doubtless a fruitful area for further research.

5.1 Optimal Changes in a Principal-Agent Game

One of the simplest ways in which multiple decision makers can be incorporated into the present problem is via a principal-agent framework.¹¹ In particular, let us assume that associated with the organization, ζ , there exists a Principal,

¹¹We are indebted to David Rode for this suggestion.

whose payoffs are given by π_P , and a set of Agents, \mathbf{A} , each with individual payoff function π_a and transformation universe \mathbb{T}_a . It is assumed that (under conditions of perfect information) the Principal wishes to take ζ from \mathbf{G}_α to \mathbf{G}_ω , but is unable to do so directly. The problem facing the Principal, then, is to select the Agent (would-be Designer) whose solution to the optimal change path problem is most highly valued.

To solve this problem, we begin by denoting by $\mathcal{P}_{\pi_{\mathbb{T}}}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)$ the optimal change path from \mathbf{G}_α to \mathbf{G}_ω in the change multigraph $\mathbf{G}_{\mathbb{T}}$ under payoff function π . Then, the solution for the Principal is to find an agent $a \in \mathbf{A}$ such that one of the following is satisfied:

$$\pi_P(\mathcal{P}_{\pi_{\mathbb{T}_a}}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)) = \max_{a' \in \mathbf{A}} \pi_P(\mathcal{P}_{\pi_{\mathbb{T}_{a'}}}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)) \quad (24)$$

$$\pi_P(\mathcal{P}_{\pi_{\mathbb{T}_a}}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)) = \max_{a' \in \mathbf{A}} \pi_P(\mathcal{P}_{\pi_{\mathbb{T}_{a'}}}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)) \quad (25)$$

$$\pi_P(\mathcal{P}_{\pi_{\mathbb{T}}}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)) = \max_{a' \in \mathbf{A}} \pi_P(\mathcal{P}_{\pi_{\mathbb{T}_{a'}}}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)) \quad (26)$$

Each of these criteria corresponds to the subgame-perfect Nash equilibrium for a different potential game.¹² Equation 24 reflects a game with unenforceable commitments and differing individual capabilities. That is, the Principal cannot enforce his or her preferences on the Agent (who, it is assumed, will maximize his or her own payoffs), and each Agent has a potentially unique set of transformations which may be applied to the optimal change problem.¹³ In this case, then, the Principal examines the solutions which each agent is anticipated to produce, and selects an agent whose solution produces the highest payoff under π_P .

The second criterion, Equation 25, provides the strategy for a game with enforceable commitments and differing individual capabilities. Here, it is assumed, the Principal can credibly enforce his or her own preferences *vis à vis* the Agent; for simplicity, we assume here that this enforcement is perfect, and thus that each Agent acts as if he or she had preference function π_P . While each Agent will then *attempt* to maximize π_P , the fact that each possesses a distinct transformation universe means that not every Agent will necessarily be successful in the attempt. As before, the Principal selects an Agent whose anticipated solution maximizes π_P .

Finally, Equation 26 gives the equilibrium strategy for a Principal-Agent game in which the Principal's preferences are unenforceable, but in which all Agents have the same capabilities. Here, while each Agent has the same ability to match π_P , each may differ on the extent that they prefer to do so. Given this, the Principal selects an Agent whose personally optimal solution he or she most values.

¹²We omit the case in which commitments are enforceable and all Agents have the same capabilities, since this is merely the standard optimal change problem.

¹³We further assume that the Principal's preferences over all such potential transformations are well-formed.

With respect to all of the above, it should be noted that actual solutions to the Principal-Agent change game for particular parameter choices can be obtained straightforwardly via a combination of the appropriate strategy (Equations 24-26) and the solutions generated by Algorithm 1. In cases where a solution to the optimal path problem does not exist for one or more Agents, this should be treated as a distinct outcome and evaluated accordingly (i.e., this should be given a payoff in π_P). Note that since some agents may not prefer \mathbf{G}_ω to all other states, it is not unreasonable that such defections will occur without enforcement! Thus, this approach manages to implicitly include enforcement problems related to disagreements regarding goals, as well as those which relate strictly to plans of implementation.

5.2 Pareto-Optimal Change Walks

A second simple scenario involving multiple decision makers arises from a situation in which the (ostensibly singular) Organizational Designer is actually a committee. If multiple individuals with divergent payoffs set out to plan an organizational change, what can we say regarding their eventual solution? Without delving too deeply into the context of the bargaining game played by these actors, we can still place some bounds on their decision by noting that any change walk they select must be Pareto-optimal. In terms of the present problem, Pareto-optimality is defined as follows:

Definition 6. Given a set of actors, \mathbf{A} , each with payoff function π_a and common transformation universe \mathbb{T} , let $\mathbb{W}(\mathbf{G}_\alpha, \mathbf{G}_\omega)$ be the set of all $(\mathbf{G}_\alpha, \mathbf{G}_\omega)$ walks in $\mathbf{G}_\mathbb{T}$. Then, for two walks $\mathcal{W}_i, \mathcal{W}_j \in \mathbb{W}(\mathbf{G}_\alpha, \mathbf{G}_\omega)$, \mathcal{W}_i is said to (strongly) Pareto-dominate \mathcal{W}_j (denoted $\mathcal{W}_i \triangleright \mathcal{W}_j$) iff $\exists a \in \mathbf{A} : \pi_a(\mathcal{W}_i) > \pi_a(\mathcal{W}_j)$ and $\nexists b \in \mathbf{A} : \pi_b(\mathcal{W}_i) < \pi_b(\mathcal{W}_j)$. A walk $\mathcal{W} \in \mathbb{W}(\mathbf{G}_\alpha, \mathbf{G}_\omega)$ is said to be Pareto-optimal iff $\nexists \mathcal{W}_i \in \mathbb{W}(\mathbf{G}_\alpha, \mathbf{G}_\omega) : \mathcal{W}_i \triangleright \mathcal{W}$.

Thus, it stands to reason that any change walk which is Pareto-dominated will not be selected, since there is at least one actor who would benefit from an alternative, and none who would suffer. Given this, which walks are Pareto-optimal? While this is difficult to determine in the general case, we can identify some Pareto-optimal walks which are particularly salient. Optimal change paths, for instance, are likely candidates, as is demonstrated by the following theorem:

Theorem 3. Let $\mathbb{P}_{\pi_a}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega) \subseteq \mathbf{G}_\mathbb{T}$ be the set of optimal change paths for some actor $a \in \mathbf{A}$. Then $\exists \mathcal{P}^* \in \mathbb{P}_{\pi_a}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega) : \mathcal{P}^*$ is Pareto-optimal.

Proof. Assume that $\nexists \mathcal{P}^* \in \mathbb{P}_{\pi_a}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega) : \mathcal{P}^*$ is Pareto-optimal. Then it follows that $\exists \mathcal{W} \in \mathbb{W}(\mathbf{G}_\alpha, \mathbf{G}_\omega) : \mathcal{W} \triangleright \mathcal{P}^* \forall \mathcal{P}^* \in \mathbb{P}_{\pi_a}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)$. Further, we know that $\mathcal{W} \notin \mathbb{P}_{\pi_a}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)$, since otherwise $\mathbb{P}_{\pi_a}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)$ would have a Pareto-optimal member. However, if $\mathcal{W} \notin \mathbb{P}_{\pi_a}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)$, then by definition $\pi_a(\mathcal{W}) < \pi_a(\mathcal{P}^*) \forall \mathcal{P}^* \in \mathbb{P}_{\pi_a}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)$, and hence \mathcal{W} cannot Pareto-dominate any member of $\mathbb{P}_{\pi_a}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)$. This contradicts our initial assumption, and hence $\mathbb{P}_{\pi_a}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)$ must contain some \mathcal{P}^* which is Pareto-optimal. \square

A trivial corollary of Theorem 3 is that unique optimal change paths are necessarily Pareto-optimal. Another obvious implication of this result is a simple procedure for finding Pareto-optimal walks: we pick an actor, find $\mathbb{P}_{\pi_a}^*(\mathbf{G}_\alpha, \mathbf{G}_\omega)$ using standard methods, and check each member against the others until we find one which is not dominated. Since the set of optimal change paths is likely to be small, the computational burden of this procedure will not be appreciably greater than the optimal path problem. While this does not allow us to generate the full set of Pareto-optimal walks, then, it does at least give us a fairly easy way to identify certain key candidates. Since these paths also have the property of being optimal for at least one actor, we assume that they will be of particular interest to the Organizational Designers.

6 Conclusion

As we have shown, a decision-theoretic approach to the problem of structural change in organizations offers a number of potentially useful insights for the theory of organizations. Normatively, this perspective has allowed us to identify optimal change paths for organizations under a range of assumptions (including uncertainty in implementation times and/or payoffs), to compute relative expected costs for the maintenance of homeostasis under external perturbations, to solve simple principal-agent problems involving the selection of Designers, and to find Pareto-optimal change paths for scenarios involving multiple decision makers with conflicting interests. Descriptively, we have been able to demonstrate minimal conditions for structural homeostasis in organizations, and to provide bounds on the expected time spent in equilibrium for various perturbation regimes. More importantly, by supplying a normative model for organizational change, we have also provided a descriptive baseline from which to assess the change behavior of real organizations.

While our theorizing has covered a fair amount of ground, there are numerous important directions which here go unconsidered. We have, for instance, given somewhat short shrift to the problem of decisions under uncertainty for Designers who are not risk neutral; this is certainly an issue which should be considered more deeply in future work. Another important assumption whose relaxation is of great interest is that of destination certainty. While we have quite explicitly presumed framed our analysis in terms of a well-defined change problem, it seems probable that, in many situations, the Designer will not be certain as to which destination is optimal. The identification of paths which are optimal under destination uncertainty is then complicated by the need to value the embedded options associated with each path (i.e., the payoff associated with the opportunity to change destination during the change process). This is a worthy topic for further development. Above all, we naturally recognize that theory alone can only take us so far. Measurement of structural and transformation payoffs, techniques for the identification of minimal transformation universes, and the like are essential for the practical application of these techniques. It is our hope that, by providing a formal decision-theoretic framework in which the

problem of organizational change can be analyzed, we will contribute to further advances in the field.

7 References

- Ahuja, R. K., Magnanti, T. L., and Orlin, J. B. (1993). *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, Upper Saddle River, NJ.
- Arrow, K. J. (1974). *The Limits of Organization*. W.W. Norton, New York.
- Baligh, H., Burton, R. M., and Obel, B. (1990). Devising expert systems in organization theory: The Organizational Consultant. In Masuch, M., editor, *Organization, Management, and Expert Systems*, pages 35–57. Walter de Gruyter, Berlin.
- Blau, P. M. (1972). Interdependence and hierarchy in organizations. *Social Science Research*, 1:1–24.
- Burton, R. M. and Obel, B. (1984). *Designing Efficient Organizations: Modelling and Experimentation*. North Holland, Amsterdam.
- Carley, K. M. and Krackhardt, D. (1998). A PCANS model of structure in organization. In *1998 International Symposium on Command and Control Research and Technology*, Monterey, CA.
- Carley, K. M. and Krackhardt, D. (1999). A typology for C^2 measures. In *1999 International Symposium on Command and Control Research and Technology*, Newport, RI.
- Carley, K. M. and Svoboda, D. M. (1996). Modeling organizational adaptation as a simulated annealing process. *Sociological Methods and Research*, 25(1):138–168.
- Cook, W. J., Cunningham, W. H., Pulleyblank, W. R., and Schrijver, A. (1998). *Combinatorial Optimization*. John Wiley and Sons, New York.
- Crowston, K. (1994). Evolving novel organizational forms. In Carley, K. M. and Prietula, M. J., editors, *Computational Organizational Theory*, pages 19–38. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Dawes, R. M. (1996). Behavioral decision making and judgment. In Gilbert, D., Fiske, S., and Lindzey, G., editors, *The Handbook of Social Psychology*. McGraw-Hill, Boston, MA.
- Dijkstra, E. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271.
- Duncan, R. B. (1979). What is the right organizational form? *Organizational Dynamics*, Winter, 1979:59–79.

- Galbraith, J. (1977). *Organizational Design*. Addison-Wesley, Reading, MA.
- Glance, N. S. and Huberman, B. A. (1994). Social dilemmas and fluid organizations. In Carley, K. M. and Prietula, M. J., editors, *Computational Organizational Theory*, pages 217–239. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Hannan, M. T. and Freeman, J. (1977). The population ecology of organizations. *American Journal of Sociology*, 82(5):929–964.
- Kalos, M. H. and Whitlock, P. A. (1986). *Monte Carlo Methods, Volume I: Basics*. John Wiley and Sons, New York.
- Koza, J. R. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA.
- Krackhardt, D. (1986). The snowball effect: Turnover embedded in communication networks. *Journal of Applied Psychology*, 71:50–55.
- Levitt, R. E., Cohen, G. P., Kunz, J. C., Nass, C. I., Christiansen, T., and Jin, Y. (1994). The “Virtual Design Team”: Simulating how organization structure and information processing tools affect team performance. In Carley, K. M. and Prietula, M. J., editors, *Computational Organizational Theory*, pages 1–18. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Lin, Z. (1994). A theoretical evaluation of measures of organizational design: Interrelationship and performance predictability. In Carley, K. M. and Prietula, M. J., editors, *Computational Organizational Theory*, pages 113–159. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Lincoln, J. R., Hanada, M., and McBride, K. (1986). Organizational structures in Japanese and U.S. manufacturing. *Administrative Science Quarterly*, 31:338–364.
- Mackenzie, K. D. (1978). *Organizational Structures*. AHM, Arlington Heights, IL.
- Marschak, T. (1986). Organizational design. In Arrow, K. J. and Intriligator, M. D., editors, *Handbook of Mathematical Economics, Volume III*. North-Holland, Amsterdam.
- Perrow, C. (1970). *Organizational Analysis: A Sociological View*. Wadsworth, Belmont, CA.
- Pfeffer, J. and Salancik, G. R. (1978). *External Control of Organizations*. Harper and Row, New York.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. (1986). *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, Cambridge.

- Simon, H. A. (1973). Applying information technology to organizational design. *Public Administration Review*, 33:268–78.
- Thompson, J. D. (1967). *Organizations in Action: Social Science Bases of Administrative Theory*. McGraw Hill, New York.
- Tversky, A. and Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, 59(4–2):S251–S278.
- Williamson, O. (1975). *Markets and Hierarchies, Analysis and Antitrust Implications: A Study in the Economics of Internal Organizations*. Free Press, New York.