Running head: The computational challenges of pursuing multiple goals

The Computational Challenges of Pursuing Multiple Goals:

Network Structure of Goal Systems Predicts Human Performance

Daniel Reichman¹, Falk Lieder², David D. Bourgin³, Nimrod Talmon⁴, &
Thomas L. Griffiths⁵

Author Note

A preliminary version of our findings about the effect of tree-width on means-selection appeared in the Proceedings of the 39th Annual Meeting of the Cognitive Science Society under the title: The Structure of Goal Systems Predicts Human Performance

¹ Corresponding author. Department of Electrical Engineering and Computer Science, University of California, Berkeley, United States. Email: daniel.reichman@gmail.com Equal contribution

² Helen Wills Neuroscience Institute, University of California, Berkeley, United States. Equal contribution

³ Department of Psychology, University of California, Berkeley, United States. Equal contribution

⁴ Department of Industrial Engineering and Management, Ben-Gurion University, Beer Sheva, Israel

⁵ Department of Psychology, University of California, Berkeley, United States

ABSTRACT

Extant psychological theories attribute people's failure to achieve their goals primarily to failures of self-control, insufficient motivation, or lacking skills. We develop a complementary theory specifying conditions under which the computational complexity of making the right decisions becomes prohibitive of goal achievement regardless of skill or motivation. We support our theory by predicting human performance from factors determining the computational complexity of selecting the optimal set of means for goal achievement. Following previous theories of goal pursuit, we express the relationship between goals and means as a bipartite graph where edges between means and goals indicate which means can be used to achieve which goals. This allows us to map two computational challenges that arise in goal achievement onto two classic combinatorial optimization problems: Set Cover and Maximum Coverage. While these problems are believed to be computationally intractable on general networks, their solution can be nevertheless efficiently approximated when the structure of the network resembles a tree. Thus, our initial prediction was that people should perform better with goal systems that are more tree-like. In addition, our theory predicted that people's performance at selecting means should be a U-shaped function of the average number of goals each means is relevant to and the average number of means through which each goal could be accomplished. Here we report on six behavioral experiments which confirmed these predictions. Our results suggest that combinatorial parameters that are instrumental to algorithm design can also be useful for understanding when and why people struggle to pursue their goals effectively.

keywords: goal systems, computational complexity, bounded rationality, graph theory

1. Introduction

In their daily lives, people pursue many goals in parallel. For instance, when people are asked to list the personal projects they are currently pursuing, each person lists about 15 projects on average (Little & Gee, 2007). People's goals are often highly interdependent and hierarchically organized such that attaining a high-level goal requires attaining multiple subgoals (Austin & Vancouver, 1996; Carver & Scheier, 1998; Miller, Galanter, & Pribram, 1960). Thus, effective goal pursuit does not only hinge on just being able to achieve a single goal. Rather, it depends on one's ability to effectively manage multiple goals simultaneously (Cantor & Blanton, 1996; DeShon & Gillespie, 2005; Fishbach, Zhang, & Koo, 2009; Johnson, Chang, & Lord, 2006; Kruglanski, Shah, Fishbach, Friedman, Chun, & Sleeth-Keppler, 2002; Louro, Pieters, & Zeelenberg, 2007; Shah, Friedman, & Kruglanski, 2002).

It is well documented that people often fail to achieve their goals (Gollwitzer, 1999; Webb & Sheeran, 2006) and such failures can result in a wide range of adverse effects (Baumeister & Heatherton, 1996). Extensive work has examined possible causes for such failures and how they can be prevented. One line of research suggests that people often fail to take actions towards their goals because they are disproportionately swayed by immediate and proximal reward relative to distant future rewards of achieving a long-term goal (Ainslie & Haslam, 1992; Green, Fristoe, & Myerson, 1994; Myerson & Green, 1995). This cognitive myopia can manifest in impulsivity (Mischel, 1974) and procrastination (Steel, 2007; Ariely & Wertenbroch, 2002), both of which can be serious obstacles to goal achievement. For instance, the inability to delay gratifications and resist temptations has been linked to lower achievements in various domains (Mischel, 1974; Mischel, Shoda, & Peake,

1988). Another line of research suggests that self-control is a limited resource that may get drained in stressful circumstances leading to self-regulatory failure (Muraven & Baumeister, 2000). Other obstacles include insufficient persistence in the face of obstacles (Bandura, 1997), forgetting to initiate goal-directed actions when the opportunity arises (Gollwitzer, 1999), failing to shield goal-striving from distractions and interruptions (Gollwitzer, 1999; Shah et al., 2002), and inadequate monitoring of goal progress (Baumeister & Heatherton, 1996; Carver & Scheier, 1998; Liberman & Dar, 2009). Other theories suggest that people fail to achieve their goals because they lack critical skills (Stajkovich & Luthans, 1998) or are insufficiently motivated to invest the effort to acquire such skills (Kruglanski, Bélanger, Chen, Köpetz, Pierro, & Mannetti, 2012).

The prevalence of people's failures to act on their goals raises the question of whether such failures are primarily due to lapses of motivation or whether other factors also play a major role. One such factor might be that having many interdependent goals creates challenging computational problems that are hard to solve in a reasonable amount of time. These challenges can be captured by the *computational complexity* of the computational problems people have to solve in order to achieve their goals.

Roughly speaking, the computational complexity of an algorithmic problem is the number of basic operations that a well-defined computational system (such as a Turing machine) has to perform to solve that problem (Papadimitriou, 2004). Here, we examine the computational complexity of problems entailed by goal pursuit and investigate whether it can discern situations where people achieve their goals from situations where they fail. Our first step is to identify the computational problems

people need to solve in order to achieve their goals. To do so, we build on the established psychological theory that people's goals and the means available to pursue them are organized into a network known as their *goal system* (Kruglanski, et al., 2002). A goal system is a *graph* where each vertex is a goal or a means and an edge connecting a means to a goal indicates that the means corresponding to one endpoint of the edge can be used to achieve the goal on the other endpoint. This theory suggests that at least some problems which people need to solve to achieve their goals, such as allocating limited resources including time and money between multiple means, can be formalized as combinatorial optimization problems over graphs. We consider two specific means-selection tasks that arise from goal systems: Selecting means in order to achieve the largest possible number of goals with a limited number of means (Max-Coverage, or MC in short) or to achieve all goals with as few means as possible (Set Cover, or SC in short). These problems are known to be computationally intractable, or NP-hard in the language of computer science (Garey & Johnson, 1979; Kleinberg & Tardos, 2005)⁶ and this intractability might be part of the reason why people often fail to achieve their goals.

Based on the principle that psychological processes are often characterized by underlying *efficient* algorithms (Sanborn, Griffiths & Navarro, 2010; Griffiths, Lieder, & Goodman, 2015), our main idea is that graph-theoretic properties that are used by computer scientists to determine the difficulty of instances of MC and SC should also be instrumental in quantifying how difficult it will be for people to solve MC and SC

⁶ For a precise definition of the term “NP-hard” and more background on this concept we refer the interested reader to Goldreich, (2008), Goldreich (2010), Harel and Feldman, (2004), and van Rooij (2008). Intuitively, NP-hardness of a computational problem means that no efficient algorithm for solving the problem exists.

instances as well. Therefore, it should be possible to predict how well people can solve problems such as MC and SC given a network of goals and means from graph-theoretic properties of the underlying goal system.

Correspondingly, here we study how well structural properties of people's goal systems can predict the success or failure of their goal pursuits by focusing on the following graph-theoretic properties: The similarity of a graph to a tree and the number of edges connected to a vertex averaged across all edges in the network (the average degree). Many NP-hard optimization problems over graphs, such as networks of goals and means, often become easier when the graph structure is *tree-like* (Kleinberg & Tardos, 2005; Kloks, 1994; Robertson and Seymour, 1986) where the term *tree* is defined as a connected graph without cycles. Furthermore, there is vast evidence that many optimization problems on networks that are “far” from trees are likely to be hard to solve exactly or approximately (Clementi & Trevisan, 1999; Marx, 2007; Trevisan, 2001). Another graph-theoretic property that is related to the hardness of solving optimization problems over networks is the average degree, i.e., the average number of connections of each vertex. Namely, as the average degree of a graph increases, the computational complexity of solving⁷ the problem increases as well (Khot & Regev, 2005). Based on these connections between graph properties and algorithmic hardness, we hypothesize that goal systems that are arranged in a network that is “tree-like” will be easier for people to deal with and enable them to perform better. With respect to the average degree, we predict a negative relationship between the average degree up until a point where every means is connected to “almost” every goal where the problem should again become easier as the degree

⁷ These hardness results refer to SC and consider hardness of approximation rather than finding the optimal solution.

increases further, because a random set of means is likely to be an optimal solution. We therefore predict that human performance is a U-shaped function of the average degree of the network connecting their goals to the means available to pursue them.

We rigorously test these predictions in experiments where people are tasked to solve MC and SC instances on goal systems that are presented to them. These findings offer a fresh computational perspective on why people find it difficult to achieve their goals. To the best of our knowledge, this article is one of the first to empirically examine whether combinatorial parameters associated with NP-hard problems can be used in a meaningful way to predict how difficult (or easy) it will be for people to cope with these problems. Our experimental results align well with theoretical knowledge from computer science and combinatorial optimization and support the relevance of findings from algorithm design and computational complexity to cognitive psychology.

This paper is organized as follows: We begin by formalizing goal systems and the computational problems entailed by goal pursuit. In the subsequent two sections we derive and empirically confirm the predictions that people's performance at deciding which means to employ to achieve a set of goals critically depends on three structural properties of the goal system: Tree-likeness, the average number of means relevant to each goal (equi-finality), and the average number of goals that each means contributes towards (multi-finality). We close by summarizing our findings and suggesting directions for future research.

2. Computational problems arising from pursuing multiple goals

People often simultaneously pursue multiple goals with hierarchically nested subgoals (Little & Gee, 2007; Newell, Shaw, & Simon, 1958). In line with this observation, previous psychological research has suggested that people mentally represent their goals, subgoals, and the means available to attain them as nodes in interconnected networks known as "goal systems" (Bandura & Cervone, 1983; Shah & Kruglanski, 2000; Kruglanski et al., 2002; Fishbach, Shah, & Kruglanski, 2004; Huang & Zhang, 2013; Louro et al., 2007). In this section, we formalize goal systems and the computational problems they entail using graph theory.

2.1 Goal systems

Using a graph-theoretic formalism², we can describe a *goal system* as a bipartite graph with a vertex set (M, G) where $M = \{m_1, \dots, m_k\}$ is a set of means and $G = \{g_1, \dots, g_l\}$ is a set of goals. We will generally assume that at least *one* means connected to a given goal has to be acquired to achieve the goal. As an example, consider the goal system depicted in Figure 1. The vertices at the top of the goal system illustrated in Figure 1 might correspond to a person's goals to become a prolific scientist (g_1), be a wonderful partner (g_2), become a great parent (g_3), get physically fit (g_4), and enjoy life to the fullest (g_5). Finally, each goal is assigned a nonnegative number that measures how desirable the goal is. We refer to this number as the *value* of the goal.

² We refer the reader to Appendix A for some graph-theoretical definitions that we use. We assume some basic familiarity with concepts from computer science such as the definition of running time of algorithms and NP-completeness whose presentation is beyond the scope of this work. More background on these topics can be found in Garey and Johnson (1979), Kleinberg and Tardos (2005), and van Rooij (2008).

We mention that it also makes sense to consider more complicated weighted models, where edges and means have weights and a goal is achieved if the weighted sum of chosen means attached to it exceeds a given threshold. To simplify the analysis and the experiments we will not deal with such weighted models here.

[INSERT FIGURE 1]

Goal systems give rise to means-selection problems where one seeks to choose a set of means that are instrumental to the ends one is trying to achieve (Kruglanski et al., 2002). Such choices are usually constrained by limited resources such as time, money, and opportunities. These resource constraints can manifest themselves in several ways. For example, it is often desirable to find the set of means that attains one's goals with the least amount of time, money, and effort possible (Shenhav et al., 2017). In other situations, people have a limited amount of time or money and aim to achieve as many of their goals as possible. The challenge then is to select a set of means that is feasible constraints and achieves as many valuable goals as possible. Here we formalize these two computational challenges in terms of well-defined computational problems which are known to be NP-hard. To illustrate these problems, suppose that you visit London for an academic conference and would like to explore the various attractions of the city. Getting a hotel in a good and central location is one means towards achieving this goal. Another means could be to buy a public transport ticket for the entire visit. A third means could be to rent a bicycle, and a fourth could be to take one of the many cruises on the Thames. You can also take the red tourist bus. Acquiring all of these means would be too costly. So you have to find a clever combination of a small number of suitable means instead.

2.2 Computational problems

We focus on the following computational problems that arise in goal pursuit.

Maximum-Coverage. The first computational problem we consider is trying to achieve as many goals as possible with a fixed budget that limits the number of activities that one can perform, corresponding to the number of means one can utilize. The budget may arise as the available means compete for limited resources such as time, money, and effort. We formalize this challenge in terms of the Maximum-Coverage problem (MC). In the Maximum-Coverage problem we are given a bipartite graph (C, D, F) , where C is one side of the bipartition, D is the second part, and F is a set of edges where each edge connects a vertex from C to a vertex in D . We are also given a nonnegative integer $k \leq |C|$. We seek to find a set $C' \subseteq C$, of cardinality at most k , maximizing the total value of vertices in D covered by vertices in C' (a node $d \in D$ is covered by a node $c \in C'$ if the edge (c, d) belongs to F). Recall that we assume that covering a goal by a *single* means is sufficient for that goal to be achieved.

Set Cover. The second problem that we study is achieving a given set of goals as efficiently as possible by selecting a set of means of minimal size that will accomplish all goals. We formalize this challenge in terms of the Set Cover problem (SC). In the Set Cover problem we are given a bipartite graph $G = (A, B, E)$ where A, B are the sides of the bipartition and E is the set of edges connecting vertices in A to vertices in B . Our goal is to cover all vertices in B by a set $A' \subseteq A$ of minimal cardinality.

The two optimization problems defined above (i.e., MC and SC) are in alignment with previous theoretical work on goal pursuit. The assumption that one

seeks to maximize the total value of attained goals is in line with expectancy-value theory originating from utility theory in economics as well as goal system theory (Kruglanski et al., 2002; Kruglanski & Kopetz, 2009; 2010; Vroom, 1964). That managing multiple goals has to occur under limited resources is widely acknowledged as well (Baumeister, Bratslavsky, Muraven, & Tice, 1998; Kahaneman, 1973; Kruglanski et al., 2002). Finally, representing goals and means as graphs and assuming interconnectedness between goals and means appear either implicitly or explicitly in several settings (Carver & Scheier, 1998; Kruglanski et al., 2002; Louro, et al., 2007). It is worth noting that, while it is suggested that people seek to *maximize* the value of goals they attain (Kruglanski & Kopetz, 2009; 2010), hardly any work has addressed the question of *how* exactly such maximization can be done in a reasonable amount of time.

2.3 Using NP-Hardness to predict human behavior

In this section we deal with few additional technical issues related to using NP-hardness to predict behavior. The reader wishing to avoid these technicalities can skip this section with no loss of continuity.

Roughly speaking, a computational problem is said to be *in NP* if there is an efficient (polynomial-time) algorithm that certifies feasibility of a proposed solution. For example, there is a simple efficient algorithm that can attest to the fact that a collection of subsets is indeed a set cover (i.e., indeed covers all means) for an SC instance, and thus, the computational problem of determining whether an SC instance has a solution of size t (for an integer t) is in NP. Informally, a problem is NP-hard if an efficient algorithm for it entails that every problem in NP has an efficient algorithm. It is widely agreed among computer scientists that NP-hard problems are

very likely to be intractable and not admit a polynomial-time algorithm (e.g., Garey & Johnson, 1979; Kleinberg & Tardos, 2005; Papadimitriou, 2004). Since it is also agreed that the non-existence of polynomial-time algorithms entails computational intractability, computer scientists usually believe that NP-hard problems are computationally intractable, in the sense that there no efficient algorithms for them exist (notice that the claim is not merely that such efficient algorithms are currently not known, but indeed that such algorithms do not exist thus cannot be found). Both MC and SC are known to be NP-hard (Garey & Johnson, 1979).

The idea that NP-hardness constrains people's ability to effectively deal with goal systems appears intuitively plausible and is in line with theoretical reasoning in psychology and cognitive science. Indeed, the NP-hardness of many computational tasks that arise in perception, decision making, problem solving, reasoning, and other domains has been recognized by psychologists (e.g., Bossaerts & Murawski, 2017; Gigerenzer, 2008; MacGregor & Ormerod, 1996; van Rooij, 2008; van Rooij, Stege, & Kadlec, 2005; van Rooij & Wareham, 2012; van Rooij, Wright & Wareham, 2012; Thagard & Verbeurgt, 1998), and the general consensus is that people are unable to find optimal solutions to NP-hard problems. Nevertheless, using NP-hardness results to explain the difficulties people may experience in dealing with goal systems is not as straightforward as it may appear. Indeed, there are several limitations of notions from complexity theory (such as NP-hardness) that may hinder their use towards predicting performance of both minds and machines. For example, complexity theory relies on asymptotic analysis that assumes the size of the instances are sufficiently large and the complexity of an algorithmic problem is measured in terms of its performance on the worst possible input. For a thorough discussion of these

limitations and possible explanations why complexity theory might be of relevance to cognitive processes we refer the reader to van Rooij (2008). These caveats of complexity theory have led to increasing reliance on empirical examinations in computer science when determining factors that influence how hard an instance of an NP-hard problem is likely to be to solve (Leyton Brown et al., 2009; Leyton-Brown, Hoos, Hutter & Xu, 2014). We therefore believe that such empirical approach is also called for when evaluating how people cope with NP-hard problems and whether theoretical results from complexity theory can be leveraged to predict human behavior.

Empirical examination of the connection between NP-hardness and human performance is also not straightforward. One difficulty is that NP-hardness is a property of an algorithmic problem, not the instances of that problem. While some instances of an NP-hard problem may be difficult for people, other instances may be easy for them. For example, the observation that MC is NP-hard is of little use if we want to determine whether two MC instances (whose underlying goal systems have the same number of goals and means) are likely to differ significantly in the number of computational steps that are needed to solve them. The difficulty in drawing a conclusion from NP-hardness to how challenging specific instances are likely to be seriously limits the utility of NP-hardness and related measures of computational complexity for predicting human performance. To overcome this difficulty, we draw on the rich literature within theoretical computer science, artificial intelligence, and combinatorial optimization that identifies structural properties of individual⁸ instances

⁸ Of course, for a single instance there is a trivial algorithm that solves the optimization problem in a constant amount of time. For example, for an MC instance the algorithm that simply prints the maximal number of goals that can be achieved solves the problem. However, having such an instance-specific algorithm in people's heads appears

of NP-hard problems that can render them easy to solve (Kleinberg & Tardos, 2005, Leyton-Brown, Nudelman, & Shoham, 2009; Hutter, Xu, Hoos & Leyton-Brown, 2014). We use this literature to derive concrete predictions about people's performance at selecting the right means to achieve goals effectively, and test them in dedicated experiments. We observe that there appears to be much less empirical evidence for the utility of computational complexity theory and algorithm design in explaining human performance and behavior (but see Enemark, McCubbins, Paturi & Weller, 2011). While there have been empirical efforts to examine which strategies people use in solving NP-hard problems (e.g., Carruthers, Masson & Stege, 2012; Dworkin & Kearns, 2015; MacGregor, Ormerod & Chronicle, 2000), these efforts do not aim at validating or refuting concrete predictions regarding what features of such problems make them hard for people.

In summary: Both MC and SC are NP-hard. The hardness of these problems gives a first indication of why people might find it so difficult to juggle multiple goals at the same time. To devise testable hypotheses as to which instances will be harder for people, it is necessary to define performance metrics that quantify how well people perform in MC and SC as well as predictors that can be used to predict human performance.

2.4 Performance metrics

People's performance at solving optimization problems can be measured in multiple ways. Our first performance measure simply quantifies the ability to find the optimal solution. This is a binary measure that equals 1 if the person finds an optimal

psychologically implausible and indeed our empirical results point against the possibility of people using such an algorithm.

solution and 0 otherwise. For SC, this means finding a cover of minimal size and for MC, this means finding a set of k means maximizing the covered goals.

Our second measure is the approximation ratio achieved on the instance of the problem. For an MC instance, we define OPT_{MC} to be the set of means of cardinality k that cover the goals of maximal total value (where the maximum is taken over all sets of means of size k). Let SOL_{MC} be the value of goals covered by a set of means chosen by a person that solved an MC instance. In this case, the approximation ratio is defined to be SOL_{MC} / OPT_{MC} . For SC we adopt a similar definition tailored for minimization problems: Let OPT_{SC} be the size of a Set Cover of minimal cardinality and let SOL_{SC} be the size of a solution returned by a person solving SC. Then, the approximation ratio is defined as OPT_{SC} / SOL_{SC} . Observe that the approximation ratio is always nonnegative and at most 1 and that the closer the approximation ratio to one, the better the performance is considered.

For both MC and SC we shall also refer to the approximation ratio as the *solution quality*. The subscripts will be omitted when the optimization problem will be clear from the context. The approximation ratio is widely used to analyze the quality of approximation algorithms for NP-hard problems (Vazirani, 2003).

In the remainder of this paper, we apply computational complexity theory to predict human performance in means-selection problems (specifically, MC and SC) and evaluate these predictions empirically. Our approach draws inspiration from empirical hardness models in computer science (e.g., Hoos, Hutter, Leyton-Brown, & Xu, 2014; Leyton-Brown et al., 2009) which leverage tools from machine learning and statistics to predict the performance of algorithms. Concretely, the following two sections evaluate whether graph-theoretic measures of tree-likeness as well as multi-

and equi-finality can predict the performance of the cognitive strategies people use to decide which means to employ to achieve a given set of goals.

3. The tree-likeness of a goal system predicts performance

In this section, we derive, test and confirm the complexity-theoretic prediction that people should perform better at means-selection when the goal system is tree-like. We start by introducing suitable graph-theoretic measures of tree-likeness.

3.1 Measures of tree-likeness

Recall that a *tree* is a *connected* network that has *no cycles*. Intuitively a tree has the minimal number of edges that ensure connectivity: the lack of cycles imply that removing even a single edge will result with a disconnected graph. Here we describe several measures of tree-likeness, namely Treewidth, combinatorial expansion, spectral expansion, and others. Each of these measures is a well-known measure quantifying different aspects of measuring closeness to a tree of a given graph.

3.1.1 Treewidth

When restricted to trees, many NP-hard problems can be solved by efficient polynomial-time algorithms, such as divide-and-conquer methods and greedy algorithms (Cormen, Leiserson, Rivest & Stein, 2001). Treewidth is a combinatorial parameter, which is always a nonnegative integer, that can be associated with any given graph. Intuitively, a smaller treewidth value implies that the nodes and edges of the network can be arranged in a way that resembles a tree (e.g., Kloks, 1994; Robertson and Seymour, 1986). For an exact definition of treewidth see Appendix A.

It is known that a connected graph has treewidth 1 if and only if it is a tree (e.g., Kleinberg & Tardos, 2005) and that, for every $d \geq 3$, there are d -regular graphs (a d -regular graph is a graph where each vertex has exactly d arcs attached to it) that have treewidth of size cn where n is the number of vertices and c is some fixed positive constant (Perarnau and Serra, 2014).

3.1.2 Combinatorial expansion

Intuitively, a given set S is *expanding* if “most” of the edges incident to S touch vertices outside S . See Figure 4 for an illustration. The edge (vertex) expansion quantify the degree of edge (vertex) expansion of the graph by taking the “worst” expanding subset of vertices that is not too large. For a precise definition see Appendix A.

Trees of bounded degree have the smallest possible expansion among connected graphs (of order $1/|V|$), hence large expansion suggests that the graph is dissimilar to a tree.

3.1.3 Spectral expansion

This measure of expansion is a real number that is positively correlated with tree-likeness (larger value implies greater similarity to a tree). This measure is based on eigenvalues of a proper matrix and efficient ways to calculate these eigenvalues are known. For a precise definition see Appendix A.

While treewidth, expansion and the spectral gap are likely to be positively correlated, there are differences between these predictors and thus we examined all of them. For example, the various expansion measures can be drastically affected by a single “bottleneck” (e.g., a graph composed of two connected components with no

edges between them). This suggests a weakness of these predictors, as a graph might be composed of two disconnected “hard” graphs and nonetheless have low expansion. Treewidth is more robust, as a union of two graphs with high treewidth that are disconnected (or connected by a small number of edges) results in a graph of large treewidth. However, the spectral gap admits an efficient polynomial algorithm as opposed to the other predictors which are NP-hard making it more practical to use.

3.1.4 Additional predictors

A feedback vertex set (FVS) is a subset of vertices whose removal from given graph results in a forest. The size of a minimal feedback vertex set is an alternative measure to the similarity of a graph to a tree. Hence we used this feature as well⁹. Previous empirical hardness models have found additional features of graphs to be useful: The diameter, the average eccentricity, and the average path length (Leyton-Brown, Nudelman, and Shoham 2009). We thus included these features in our analysis as well. The diameter of a graph is the longest distance between two vertices in the graph (where the distance is the number of edges in a shortest path connecting two vertices; all graphs in this work are connected). The eccentricity $\varepsilon(v)$ of a vertex v in an undirected graph $G = (V, E)$ is the maximal distance of a vertex in V (different from v) from v . The average eccentricity (AvgEcc) is defined as $\frac{1}{|V|} \sum_{v \in V} \varepsilon(v)$. The average path length (AvgPath) is $\frac{2}{|V|(|V| - 1)} \sum_{u, v \in V, u \neq v} d(u, v)$ where $d(u, v)$ is the distance (shortest length) between u and v (and every pair of vertices is counted in the summation once).

We calculated FVS using an implementation based on (Iwata 2016; Wahlström, 2014) (available at <https://github.com/wata-orz/fvs>).⁹

3.2 *The connection between treelikeness and efficient algorithms*

Treewidth has found numerous algorithmic applications (Robertson & Seymour, 1986) and also proven instrumental to solving inference problems on graphical models (Kask, Gelfand, Otten & Dechter, 2011). The crucial observation is that for many algorithmic problems, an n -vertex graph of treewidth w admits an exact divide-and-conquer based algorithm that solves the problem in time proportional to $2^w n$ (Bodlaender, 1988). Both MC and SC are no exception: An algorithm with running time $O(2^w n)$ is known for n -vertex bipartite graphs whose underlying factor graph has treewidth w (Alber & Niedermeier, 2002). On the other hand, graphs with large treewidth appear to be difficult to solve: For some problems, the exponential dependency of the running time on the treewidth w (for algorithms which find the optimal solution) is inherent under certain assumptions (Marx, 2007). As worst-case instances of MC and SC are hard even to approximate (Feige, 1998), and as hard instances often have large treewidth (Clementi & Trevisan, 1999; Trevisan, 2001), it is likely that it will be not only hard to solve instances exactly, but also to find an approximate solution on instances of large treewidth.

Intuitively, a tree-like goal system is one in which goals can be pursued nearly independently of each other making the selection of means a task that can be decided locally which results in an easier problem. In contrast, in goal systems that are *not* tree-like there is a stronger connection between goals and the means chosen to satisfy them making means-selection a global problem, which is much more difficult. As noted above, these intuitions can be made precise: Graphs that are more tree-like admit more efficient algorithms for solving means-selection problems.

Current algorithms that compute tree-decompositions are quite complicated and time consuming. Therefore, it seems unlikely that people would use them to solve SC and MC problems. But treewidth might nevertheless affect the performance of people's heuristics: The similarity between low treewidth graphs and trees might make the kinds of algorithms that people might use, such as greedy methods, much more effective. Similar reasoning applies to our expansion measures. For example, NP-hard problems over graph families which are guaranteed to have sub-linear edge or vertex separators (e.g., planar graphs, graphs excluding a fixed minor) are known to have better approximation algorithms compared to arbitrary instances (e.g., Baker, 1994; Klein, Plotkin, & Rao, 1993). Summarizing the discussion above, we derive the following hypothesis:

Hypothesis I. Treewidth, vertex-expansion, edge-expansion, and the spectral gap of G , $d - \lambda_2$, should be negatively correlated with the solution quality and probability that a person will find the optimal solution to the SC and MC problems.

3.3 Design of treewidth experiments

To test how our predictors can capture human performance in MC and SC, we conducted a series of behavioral experiments. To test whether human performance in means-selection deteriorates with increasing treewidth (Hypothesis I), we used a between-subjects design, where each participant was assigned randomly to one of twenty graphs. In the Experiment 1, participants were asked to solve the SC problem, while in the Experiment 2 participants were asked to solve the MC problem. In each case, the problem was graphically represented as a bipartite graph. To avoid potential

influence of the problem size on the performance in these tasks we used graphs with an identical number of vertices, specifically exactly 48 vertices, and edges across all experiments. In particular, all graphs were 4-regular and thus had exactly 96 edges. The 24 vertices at the bottom represent the available means (activities A-Z) and 24 vertices at the top represent the goals. Each edge from a means vertex to a goal vertex implies that completing that activity is sufficient to achieve the goal. In the SC task, participants were asked to select a minimal number of activities to achieve all goals. In the MC task, participants were asked to choose five activities that achieve as many goals as possible. Goals often have semantic content and different goals differ in their values. Experiment 3 asked participants to solve a SC problem where goals are given semantic content and real values. Furthermore, a different visual display is used to eliminate visualization effects. The same family of bipartite graphs as in the first two experiments were used in Experiment 3.

Our experiments follow a common paradigm where people have to choose a set of means to achieve a set of goals coupled with a verbal description of some of the possible means or goals. While this paradigm clearly does not capture all aspects of goal pursuit in the wild (which are often significantly more complex) it is similar to previous experiments done within the goal systems framework (Köpetz, Faber, Fishbach, & Kruglanski, 2011; Huang & Zhang, 2013; Zhang et al., 2007).

All experiments reported below were approved by the institutional review board of the University of California, Berkeley under protocol number 2015-05-7551, study title “Decision Making”.

3.4 Experiment 1: human performance on Set Cover

3.4.1 Methods

3.4.1.1 Participants. We recruited 655 participants on Amazon Mechanical Turk.

Participants were paid \$1.25 and could earn a performance-dependent bonus of up to \$2.

3.4.1.2 Stimuli and procedure. we generated 20 graphs whose treewidth ranged from 5 to 14. Besides one random 4-regular bipartite graph, a 6 x 8 torus and a 4 x 12 torus, we generated 3 families of graphs. For each of these families we first generated a *base* graph of low treewidth by examining plausible candidate graphs and computing their treewidth; for two of the families we used a graph of treewidth 6, and for the third family we used a graph of treewidth 5. Then, we used two methods which take as input those base graphs, and generate graphs of varying higher treewidth. The first method operated by performing random flips: we randomly choose two edge (a, b) and (c, d) for which we verify that neither (a, c) nor (b, d) already belongs to the graph on which the flip is to be performed. Then, we delete the edges (a, b) and (c, d) and replace them by the edges (a, c) and (b, d) . By increasing the number of these random flips we could gradually increase the treewidth. The second method identified a matching of size 8 in the graph and reconnecting the matched vertices by a random matching. Again, by repeating this process we could gradually increase the treewidth. Each participant was randomly assigned to one of 20 conditions that differed only in the graph structure of the SC problem participants were asked to solve. After consenting to participate, participants read a cover story about a person trying to choose which set of activities (e.g., volunteer to improve the company's website and work out at the gym) they should perform in order to achieve all their goals (e.g., earn more money, improve relationship with boss, get fit, etc.) with as few activities as possible because their time is limited. The story highlighted that some activities achieve multiple goals at the same time. Next, participants completed a simple

practice trial involving only two means and two goals. Once they had solved this task successfully, participants proceeded to the Set Cover problem they had been assigned to (see Figure 5).

[INSERT FIGURE 5]

When participants hovered with their mouse over an activity, the interface highlighted the goals this activity would achieve and the corresponding edges of the graph in green. Goals that had already been achieved and activities that had already been selected were highlighted by checkmarks, and the number of selected activities was shown at the bottom of the screen. Participants were asked to help a person achieve their 24 goals with as few activities as possible, and they were motivated by the prospect of earning a financial bonus of \$2 for achieving all goals as efficiently as possible. After submitting their solution, participants completed an exit questionnaire asking them for basic demographic information (age, gender, and primary language), and four 9-point Likert scales (anchors: “not at all”, “somewhat”, and “extremely”) measuring the perceived difficulty of the task, motivation to achieve all goals, motivation to find the optimal solution, and their motivation to finish the task as quickly as possible. For more information on the interface, we refer the reader to the Supplementary Material.

[INSERT FIGURE 6]

3.4.2 Results

We excluded 30 participants (4.6%) because their responses did not achieve all goals, suggesting that they did not follow the instructions.

We found that treewidth alone explained 44.90% of the variance in the frequency with which people found the optimal solution across the 20 graphs ($F(1, 18) = 14.20$; $p = 0.0012$): As we increased the treewidth of the graph the percentage of participants who discovered the optimal solution decreased significantly ($\rho = 0.59$; $p = 0.0058$) from more than 90% on the graph with treewidth 5 to only about 30% on the graph with treewidth 14 (Figure 6a). We found that the average solution quality was negatively correlated with treewidth ($\rho = 0.44$; $p = 0.0525$) suggesting that our participants achieved fewer goals for goal systems with higher treewidth (Figure 6b). Treewidth explained 17.59% of the variance in the median response time across problems ($F(1, 18) = 3.86$; $p = 0.0650$): The median amount of time people took to solve the problems tended to increase with treewidth ($\rho = 0.3426$) but this effect was not statistically significant ($p = 0.1393$), and when we restricted this analysis to correct solutions the correlation was $\rho = 0.3825$ ($p = 0.1297$). Perceived difficulty also tended to increase with treewidth ($p = 0.37$) but this correlation was not statistically significant ($p = 0.1062$). Our participants were highly motivated to find the optimal solution (average rating 7.91 ± 0.06 out of 9). Thus, it appears unlikely that their motivation was a bottleneck to their performance. Furthermore, motivation appeared to be unaffected by treewidth ($\rho = 0.23$; $p = 0.34$). Thus, the observed differences in performance appear to result from the inherent difficulties of the means-selection problems posed by different goal systems.

Of the additional predictors evaluated, we found that graph diameter, average shortest path, and average graph eccentricity all were significantly positively

correlated with the frequency of optimal solutions identified by our participants (graph diameter: $\rho = 0.5691$; $p = 0.0088$, average shortest path: $\rho = 0.6516$; $p = 0.0019$, average eccentricity: $\rho = 0.6265$; $p = 0.0031$). In addition, the spectral expansion (measured as $d-\lambda_2$) and vertex and edge expansions showed significant negative correlations with the frequency of optimal solutions (vertex expansion: $\rho = 0.5836$; $p = 0.0069$, edge expansion: $\rho = 0.4552$; $p = 0.0437$, spectral expansion: $\rho = 0.6280$; $p = 0.0030$). We also found that the average shortest path, average graph eccentricity, spectral expansion, and the size of the graph vertex expansion were significantly correlated with the average participant solution qualities (average shortest path: $\rho = 0.4505$; $p = 0.0462$, average eccentricity: $\rho = 0.4316$; $p = 0.0574$, spectral expansion: $\rho = 0.4496$; $p = 0.0467$, vertex expansion: $\rho = 0.4636$; $p = 0.0395$). Finally, only the cardinality of the graph edge expansions exhibited a significant correlation with the median response times on the SC task ($\rho = 0.4538$; $p = 0.0445$).

None of the effects reported above can be attributed to variability in the size of the optimal solution, since it was neither correlated with the average frequency with which people found the optimal solution ($\rho = -0.0667$, $p = 0.6827$) nor with the quality of their solutions ($\rho = 0.1356$, $p = 0.4043$).

3.5 Experiment 2: human performance on Maximum Coverage

3.5.1 Methods

3.5.1.1 Participants. We recruited 545 participants on Amazon Mechanical Turk. Participants were paid \$1.25 and could earn a bonus of up to \$2. The consent form specified that participants must not have participated in the previous version of this experiment.

3.5.1.2 Stimuli and procedure. Experiment 2 was identical to Experiment 1 except for the task: Participants were now instructed to achieve as many goals as possible subject to the constraint that the person's limited time does not permit them to complete more than five activities. The 20 graphs and financial incentives were the same as in Experiment 1. The interface of Experiment 1 (see Figure 5) was modified to prevent participants from selecting more than five activities at a time. When a participant attempted to add a sixth activity they were told they would first have to remove one or more of the activities they had already selected. The cover story and survey were modified slightly to match the change in the task.

[INSERT FIGURE 7]

3.5.2 Results

We excluded 23 participants (4.2%) because they had selected fewer than five means.

The frequency with which people found the optimal solution decreased significantly with treewidth ($\rho = 0.4828$; $p = 0.0311$; Figure 7a). We found that treewidth alone explained 20.41% of the variance in the frequency with which people found the optimal solution across the 20 graphs ($F(1, 18) = 4.62$; $p = 0.0455$). We found that treewidth explained 25.25% of the variance in solution quality ($F(1, 18) = 6.08$; $p = 0.0240$) which significantly deteriorated as treewidth increased ($\rho = 0.4972$; $p = 0.0257$; Figure 7b). The median amount of time people took to solve the problems did not increase significantly with treewidth ($\rho = 0.25$; $p = 0.28$) and treewidth explained only 0.6% of our participants' median response times ($F(1, 18) = 0.10$; $p = 0.76$). When we restricted the analysis to the time taken by optimal solutions, the

relationship was still not statistically significant ($\rho = 0.2994$; $p = 0.1998$; $F(1, 18) = 0.10$; $p = 0.76$). Finally, treewidth explained only 8.8% of the variance in the perceived problem difficulty across the 20 graphs ($F(1, 18) = 1.74$; $p = 0.20$), and the correlation between treewidth and perceived difficulty was not statistically significant ($\rho = 0.26$; $p = 0.26$). Our participants were highly motivated to find the optimal solution (average rating 8.11 ± 0.06 out of 9). Thus, it appears unlikely that their motivation was a bottleneck to their performance. Furthermore, motivation appeared to be unaffected by treewidth ($\rho = 0.03$; $p = 0.91$). Thus, the observed differences in performance appear to result from the inherent difficulties of the means-selection problems posed by different goal systems.

In addition, we found that both the size of the graph edge expansion and the graph spectral expansion (measured again as $d\text{-}\lambda_2$) were significantly negatively correlated with the frequency of optimal solutions (edge expansion: $\rho = 0.4802$; $p = 0.0321$, spectral expansion: $\rho = 0.4782$; $p = 0.0330$), while the average shortest path and average graph eccentricity showed a significant positive relationship (average shortest path: $\rho = 0.4912$; $p = 0.0279$, average eccentricity: $\rho = 0.4391$; $p = 0.0528$). In contrast, only the size of the graph vertex and edge expansions showed a significant correlation with the average solution quality (vertex expansion: $\rho = 0.4431$; $p = 0.0504$, edge expansion: $\rho = 0.4832$; $p = 0.0309$), suggesting that, in general, graph treewidth and combinatorial expansions may be more robust predictors of human performance on the MC problem. None of the metrics surveyed were significantly correlated with median participant response times.

3.6 Experiment 3: A more concrete Set Cover task

While Experiments 1 and 2 capture some of the computational challenges of goal achievement the tasks were relatively abstract. Experiment 3 addresses this limitation by assigning semantic labels to the 24 goals. These labels were common new-year resolutions such as “get in shape” and “earn more money”. We also used a different interface to minimize the effect of visualization that arise from graph drawings in the first two experiments.

[INSERT FIGURE 8]

3.6.1 *Methods*

3.6.1.1 Participants. We recruited 600 participants on Amazon Mechanical Turk.

3.6.1.2 Stimuli and procedure. Participants were paid \$0.38 for about 5min of work plus a performance-dependent bonus of \$0.50 if they found an optimal solution. Each participant was randomly assigned to one of the twenty graph structures used in Experiments 1 and 2. For each graph, the order in which the means were listed and the order in which the goals were listed was randomized between participants. The participants’ task was to achieve all goals with as few means as possible. The graphical interface of the task was changed to reduce visual clutter. Instead of drawing edges between means and goals, the goals achieved by each means were listed next to it (see Figure 8). The cover story was similar to the one used in Experiment 1. The consent form required that participants had not participated in any of our previous goal pursuit experiments.

[INSERT FIGURE 9]

3.6.2 *Results*

As shown in Figure 9, Experiment 3 replicated the effects of treewidth on the solution frequency and the solution quality in a more realistic setting: Increasing the treewidth of the goal system significantly decreased both solution frequency ($\rho = -0.756, p = 0.0001$) and solution quality ($\rho = -0.60, p = 0.004$). Furthermore, we found that the magnitude of the graph spectral expansion, cardinalities of the graph edge and vertex expansions, average eccentricity, average shortest path, and graph diameter were all significantly correlated with the frequency with which human participants identified the optimal solution (average eccentricity: $\rho = 0.583; p = 0.007$, average shortest path: $\rho = 0.651; p = 0.002$, graph diameter: $\rho = 0.525; p = 0.017$, spectral expansion: $\rho = 0.708; p = 0.0005$, edge expansion: $\rho = 0.7; p = 0.0006$, and vertex expansion: $\rho = 0.7303; p = 0.0003$).

Similarly, the magnitude of the graph spectral expansion, cardinalities of the graph edge and vertex expansions, average eccentricity, and average shortest path were all significantly correlated with the average solution quality of human responses (average eccentricity: $\rho = 0.4872; p = 0.0293$, average shortest path: $\rho = 0.5243; p = 0.0176$, spectral expansion: $\rho = 0.5308; p = 0.0160$, edge expansion: $\rho = 0.5670; p = 0.0091$, and vertex expansion: $\rho = 0.6047; p = 0.0047$). None of the features were significantly correlated with median participant response times.

Participants rated their motivation to find a solution that achieves all goals with the minimal number of means as 7.38 on a 9 point scale, their motivation to finish the task as quickly as possible and move on as 4.35, and the difficulty of the task as 5.67. The size of the optimal solution was positively correlated with the average solution quality ($\rho = 0.46, p = 0.0387$) but not with the average solution frequency ($\rho = 0.25, p = 0.29$).

3.7 Discussion of treewidth experiments

In three experiments, we have demonstrated that treewidth is a reliable predictor of human performance in means-selection problems. Specifically, treewidth had a negative effect on both solution quality and the probability of finding the optimal solution. This negative effect emerged on different computational tasks (Set Cover and Max-Coverage) and in different visual displays. To conclude, these experiments support the hypothesis that larger treewidth of goal systems has a reliable and robust negative effect on human performance. One additional reason why people's performance deteriorates with increasing treewidth could be that the means-selection problems entailed by treelike goal systems require less working memory to solve than those entailed by goal systems with high treewidth.

4. The average degree of a goal system predicts performance

Previous research has postulated that the selection of means by people is affected by how many goals they are connected to and the number of means connected to the corresponding goal (Kruglanski, et al., 2002). Viewing means and goals as nodes in a graph allows us to connect this idea to computational complexity theory in a concrete way.

4.1 Graph-theoretic predictors: average degree of goals and means

The degree of a means vertex is simply the number of goal-vertices that are connected to it. Similarly, the degree of a goal vertex is the number of means-vertices connected to it. When studying the effect of the degree on performance, we only consider instances where all means vertices have the same degree or all goal vertices have an identical degree. We impose the regularity restriction to enhance our ability to

isolate the effect of the degree on the performance of people in MC and SC. We use d_{means} to denote the degree of means-vertices and use d_{goal} for goal-vertices. The degree d_{means} is a basic combinatorial parameter of bipartite graphs (and goal systems in particular) that is related to the performance of various algorithms for SC. For example, it is known that a simple greedy algorithm finds a number of means that is larger by a multiplicative factor of at most $H(d_{means})$ than the optimal number of means needed to cover all goals (Cormen, Leiserson, Rivest & Stein, 2001, and the references therein; Lovasz, 1975). Here, $H(r)$ is the Harmonic number: $H(r)=1+1/2+\dots+1/r$. Furthermore, there is some evidence that, as the degree gets larger, the approximation ratio achievable for SC (as far as polynomial-time algorithms are concerned) deteriorates as a function of d_{means} (Feige, 1998; Saket and Sviridenko, 2012).

A similar phenomenon holds for d_{goal} . A simple algorithm (iteratively picking all means covering a given goal, adding them to the cover, and deleting all covered goals, until all goals are covered) is known to achieve an approximation ratio of $1/d_{goal}$. Furthermore, hardness of approximation results indicates that achieving a significantly larger approximation ratio than $1/d_{goal}$ is likely to be intractable (Dinur, Guruswami, Khot & Regev, 2005; Khot & Regev, 2005).

The parameters d_{means} and d_{goals} also arise in goal system theory (Kruglanski et al., 2002). For example, d_{means} captures a normative element of a given means as it quantifies to what extent this means is instrumental to the achievement of several goals. That means achieve multiple goals has been referred to as the *multifinality* property of means. Furthermore, it has been observed that people are sensitive to the degree of means, and that this degree can influence how valuable a means is

perceived to the achievement of the goals adjacent to it (Zhang, Fishbach, & Kruglanski, 2007). The parameter d_{goal} is of relevance as well: A larger d_{goal} implies that there are more ways to attain the goal associated with a goal vertex. That a goal can be pursued by different means allowing for alternative ways to achieve the goal is often referred to as the *equifinality* property of goals.

How do d_{goal} and d_{means} influence people's performance on MC and SC? To simplify notation we use the notation of d for both these parameters in this paragraph. On the one hand, the more goals are covered by each means, the fewer means it should take to cover all goals (similar reasoning applies when considering the degree of goal vertices). Therefore, increasing d should be beneficial for goal-attainment as one can achieve more goals using less means. Furthermore, when d is large relative to the number of goals, the fraction of subsets of size k that are optimal solutions approaches 1. This increases the likelihood that a randomly generated solution will be optimal which appears to make the problem easier (Enemark et al., 2011). On the other hand, psychological and theoretical considerations suggest that cognitive and computational complexity increase with d : Increasing the *multifinality* (or *equifinality*) of means (goals) might increase people's cognitive load, as more goals per means should be considered when deciding whether to select means. In addition, as previously noted, there is theoretical evidence that instances of SC with larger degree are harder to approximate. Furthermore, from an approximation algorithm perspective, the decrease of the size of the solution makes the problem harder to approximate, as for minimization problems larger lower bounds on the optimal solution often translate to better approximation guarantees.

The above reasoning together with considerations of the limited size of our goal systems lead us to hypothesize that d will be negatively associated with the quality solution achieved by people as long as d is *sufficiently small* compared to d . When d approaches the number of goals (means) we expect to see a positive relationship between d and the solution quality. The average degree of a goal system corresponds to the average number of means connected to each goal (equifinality) and the number of goals that are connected to each means (multifinality). Indeed, in one extreme case, where every goal is connected to only a single means ($d = 1$), then it seems to be very easy to choose the right means for each goal. Then, adding a few connections would make the problem more challenging. This intuition is supported by both algorithmic and intractability results. In the other extreme case, where every goal is connected to every means and vice versa the problem becomes trivial again. We formalize this prediction below:

Hypothesis II: There will be a curvilinear U-shaped relationship between d and the quality of solution achieved by people in SC.

A natural question is whether the same U shape should hold for the fraction of people solving SC exactly. We are not aware of theoretical results regarding the hardness of solving SC instances *exactly* that differentiate between “large” and “small” d , though as d increases, similar reasoning to Hypothesis II, suggests that exact solution of the SC should be easier. Hence, we did not have a clear hypothesis about the effect of d on the ability to find the optimal solution for SC.

The effect of the number of edges in a network on the ability of people to solve NP-hard problems was studied by Enemark et al. (2011) using graph coloring problems (in these problems, the vertices of a given graph are to be colored by few

colors while it is forbidden to color adjacent vertices with the same color). They found that, the higher the probability a random solution is an optimal solution, the easier it was for people to solve the problem they encountered. The main measure used by Enemark et al. (2011) was random solution likelihood (RSL) which is the logarithm of the probability a random solution is an optimal solution in an NP-hard coloring task. We reasoned that, similarly to Enemark et al. (2011), RSL might explain the influence of the density of the network on people's performance in SC. Hence, we used this measure as a predictor for people's performance as well.

4.2 Experimental tests of hypothesis II

To test Hypothesis II, we conducted three experiments that examine how multifinality (i.e., the number of goals connected to each mean, d_{means}) and equifinality (i.e., the number of means connected to each goal, d_{goal}) affect people's ability to select the optimal set of means for a given set of goals. In these experiments, participants were asked to solve SC on graphs where either d_{goal} or d_{means} were systematically varied. In the multifinality experiments we examined how varying d_{means} influences performance, while in the equifinality experiments, we studied how d_{goal} influences performance. Figure 10 shows an example of the sets of graphs used in the experiment. All experiments reported below were approved by the institutional review board of the University of California, Berkeley under protocol number 2015-05-7551, study title "Decision Making".

[INSERT FIGURE 10]

4.3 Experiment 4: Multifinality and human performance in Set Cover

4.3.1 Methods

4.3.1.1 Participants. We recruited 50 participants on Amazon Mechanical Turk. Participants were paid \$3.25 and could earn a bonus of up to \$4. The median completion time of the experiment was 34.8 minutes, and the average bonus was \$0.83.

4.3.1.2 Stimuli and procedure. Each participant solved 8 means-selection problems presented using the same interface from Experiment 1 (see Figure 5). The experimental design was a within-subjects design with 8 conditions where the number of goals achieved by each means was 2, 3, 4, 5, 6, 8, 10, and 12 respectively. The order of these eight conditions was randomized between subjects. Each graph was generated by connecting every means-vertex to d randomly chosen goal-vertices.¹⁰ Furthermore, we randomized other properties of the graphs across participants by selecting the graph for each condition from a collection of 20 randomly generated graphs with the required degree. On each trial, participants were required to achieve all goals in order to submit their solution.

The consent form specified that participants must not have participated in the previous version of this experiment. Participants completed one practice trial before solving the eight problems. We used a similar visual display to the one used in Experiments 1 and 2. For each problem, participants were incentivized by a bonus of \$0.50 for solving it optimally.

[INSERT FIGURE 11]

4.3.2 Results

¹⁰ Observe that this sampling process did not guarantee that every goal vertex had the same degree, and in general the degrees of the goal vertices differed.

Participants perceived the problems as difficult (average self-reported difficulty: 6.18 ± 1.70 out of 9) but were highly motivated to solve them rating their motivation to find the optimal solution as 7.92 ± 1.1 out of 9 and their motivation to finish the tasks as quickly as possible as only 3.78 ± 2.24 out of 9.

Figure 11b shows the average quality of participants' solutions as a function of multifinality. The average solution quality differed significantly between the five conditions according to a mixed-effects ANOVA ($F(7,384) = 2.72, p = 0.0091$). To test our hypothesis that average solution quality is a U-shaped function of the number of goals per mean, we fitted a general linear model with the regressors d and d^2 . We found that, while d had a significant negative effect on solution quality ($\beta_d = -0.08$, $t(397) = -8.7, p < 10^{-15}$), d^2 had a significant positive effect on the average solution quality ($\beta_{d^2} = 0.005$, $t(397) = 7.4, p = 7.8 \cdot 10^{-13}$) and the additional variance explained by d^2 above and beyond the variance explained by d was highly significant ($F(1,397) = 54.86, p = 7.8 \cdot 10^{-13}$). To test whether the non-linear effect of d was U-shaped, we estimated the location of the minimum point (Lind and Mehlum, 2010; Bebu, Seillier, Moiseiwitsch, & Mathew, 2009): If the minimum point lies at an intermediate value of d , then the effect can be called U-shaped. But if the minimum point lies outside the range of feasible d then the effect could be still monotonic. According to our data we are 99% confident that the minimum point lies between 7.8 and 9.3 ($\hat{d}_{min} = 8.36$, 99% CI: [7.8067, 9.2847]). We can therefore conclude that the average solution quality is a U-shaped function of multifinality.

Next, we analyzed the effect of multifinality on solution frequency (see Figure 11a). We found that the frequency with which people found an optimal solution differed significantly between the 8 groups ($\chi^2(7) = 24.3, p = 0.001$). Next, we

investigated whether this relationship is U-shaped. We found that, while d had a significant negative effect on solution frequency ($\beta_d = -0.79, t(397) = -4.12, p = 3.7 \cdot 10^{-5}$), d^2 had a significant positive effect on it ($\beta_{d^2} = 0.05, t(397) = 4.04, p = 5 \cdot 10^{-5}$). The additional variance explained by adding d^2 was statistically significant ($F(1,397) = 17.05, p = 4.4 \cdot 10^{-5}$). Next, we tested whether the minimum point lies inside the range of degrees we investigated according to the method proposed by Lind and Mehlum (2010). The model's estimate for the minimum point is $d_{min} = \frac{-\hat{\beta}_x}{2 \cdot \hat{\beta}_{x^2}}$. Therefore, we have to test the alternative hypothesis that $\beta_x + 2 \cdot x_L \cdot \beta_{x^2} < 0$ and $\beta_x + 2 \cdot x_H \cdot \beta_{x^2} > 0$ against the null hypothesis which says that null hypothesis that at least one of the two statements is false. Based on the results for the solution quality we test whether the minimum lies between 6 and 9. In support of this hypothesis, we found that both components of the null hypothesis can be rejected ($F_1(1,397) = 9.78, p = 0.0019, F_2(1,397) = 8.66, p = 0.0034$) for $x_L = 6$ and $x_{high} = 9$. We therefore conclude that multi-finality has a U-shaped effect on solution frequency.

Furthermore, we found that the number of times participants clicked on means decreased monotonically with multifinality (Spearman's $\rho = -0.98, p = 0.0003$) and so did the number of means selected by the participants (Spearman's $\rho = -1, p < 0.0001$). Furthermore, we also found that the median reaction time decreased significantly with multifinality ($\rho = -0.98, p = 0.0004$).

Finally, we tested whether the effect of multifinality on relative performance was mediated by the random solution likelihood using the mediation analysis methodology by Baron and Kenny (1986). We found that controlling for the random solution likelihood by including its logarithm as a covariate in the regression equation

did neither abolish nor significantly reduce the effect of d or d^2 on relative solution quality: Without controlling for the random solution likelihood the effect of the d was $\hat{\beta}_d = -0.0772$ (95% CI: $[-0.0946, -0.0597]$) and the effect of d^2 was $\hat{\beta}_{d^2} = 0.0046$ (95% CI: $[0.0034, 0.0058]$). When controlling for the random solution likelihood the corresponding effects became $\hat{\beta}_d = -0.0958$ (95% CI: $[-0.1130, -0.0786]$) and $\hat{\beta}_{d^2} = 0.0042$ (95% CI: $[0.0031, 0.0054]$). Since the confidence intervals overlap almost completely there appears to be no change in the effect of the number of goals on average solution frequency when the random solution likelihood is controlled for. The effect of multifinality on the frequency that people would find the optimal solution was also not mediated by the random solution likelihood because the random solution likelihood did not have a significant effect on the probability that people would find the optimal solution ($\hat{\beta}_{RSL} = 0.0654, t(398) = 1.58, p = 0.11$).

The size of the minimal solution was positively correlated with people's relative performance ($\rho = 0.46, p < .0001$) and with whether or not the solution was optimal ($\rho = 0.13, p = 0.01$). In particular, smaller values of an optimal Set Cover resulted with worse performance.

Most importantly, the results of this experiment confirmed our theoretical prediction that people's ability to identify an optimal set of means for achieving their goals should be a U-shaped function of multifinality (Hypothesis II). Next, we evaluate whether Hypothesis II is also true for the number of means connected to each goal (equifinality).

4.4 Experiment 5: Equifinality and human performance in Set Cover

4.4.1 Methods

4.4.1.1 Participants. We recruited 50 participants on Amazon Mechanical Turk. Participants were paid \$3.25 and could earn a bonus of up to \$4. The median completion time of the experiment was 41.2 minutes, and the average bonus was \$1.26.

4.4.1.2 Stimuli and procedure. Participants solved 8 max-coverage problems presented using the interface from Experiment 1 (see Figure 5). The experimental design was a within-subjects design with 8 conditions where the number of means connected to each goal was 2, 3, 4, 5, 6, 8, 10, and 12, respectively. The order of these eight conditions was randomized between subjects. The graphs were generated in identical fashion to the graphs in Experiment 4. A few times, we encountered a graph where there existed a goal that was not covered by any means; in such cases we deleted the graph and resampled a new graph from the same distribution until all goals were covered by at least one means. On each trial participants were required to achieve all goals in order to submit their solution. We used a similar interface to Experiment 3, where no edges between goal and means appear and goals are given semantic content.

The consent form specified that participants must not have participated in the previous version of this experiment. Participants completed one practice trial before solving the eight problems. For each problem, participants were incentivized by a bonus of \$0.50 for solving it optimally.

[INSERT FIGURE 12]

4.4.2 Results

On average, participants rated the difficulty of tasks as 5.74 ± 1.69 (means $\pm$ SD), their motivation to find the optimal solution as 7.38 ± 1.74 on a 9-point scale, and the average motivation to finish the tasks as quickly as possible was rated as only 3.48 ± 1.96 .

The average solution quality differed significantly between the eight conditions according to a mixed-effects ANOVA ($F(7,384) = 16.89, p < 10^{-15}$; see Table 1). Figure 12 shows the effect of equifinality on the participant's ability to select the optimal set of means. To test our hypothesis that average solution quality is a U-shaped function of the number of goals per mean (Figure 12b), we fitted a general linear model with the regressors d and d^2 . We found that while d had a significant negative effect on solution quality ($\beta_d = -0.04, t(397) = -3.98, p < 0.0001$), d^2 had a significant positive effect on the average solution quality ($\beta_{d^2} = 0.003, t(397) = 3.82, p = 0.0002$) and the additional variance explained by d^2 above and beyond the variance explained by d was highly significant ($F(1,397) = 14.62, p = 0.0002$). To test whether the non-linear effect of d was U-shaped, we estimated the location of the minimum point (Lind & Mehlum, 2010; Bebu, Seillier-Moiseiwitsch, & Mathew, 2009): According to our data we are 99% confident that the minimum point lies between 6.4 and 9.2 ($\hat{d}_{min} = 7.41, 99\% \text{ CI: } [6.3787, 9.1762]$). We therefore conclude that the average solution quality is a U-shaped function of equifinality.

Next, we analyzed the effect of equifinality on solution frequency (see Figure 12a). We found that the frequency with which people found an optimal solution differed significantly between the 8 groups ($\chi^2(7) = 24.4, p = 0.001$). Next, we investigated the shape of the relationship using a logistic regression model with both

d and d^2 as regressors. We found that, while d^2 had a significant positive effect on solution frequency ($\beta_{d^2} = 0.0258, t(397) = 2.18, p = 0.0291$), the additional effect of d was not statistically significant ($\beta_d = -0.25, t(397) = -1.48, p = 0.14$). The additional variance explained by adding d to a model with the regressor d^2 was not statistically significant ($F(1,397) = 2.58, p = 0.11$). This suggests that within the investigated range the effect is neither linear nor U-shaped but approximately quadratic. However, a more accurate description may be that the solution frequency was constant for $2 \leq d \leq 10$ and then increased abruptly as d was raised to 12. Indeed, pairwise χ^2 -tests showed that there were no statistically significant differences between the solution frequencies between the groups with $2 \leq d \leq 10$ ($\chi^2(6) = 3.16, p = 0.79$), but the solution frequency for $d = 12$ was significantly higher than for $d \leq 10$ ($\chi^2(1) = 21.51, p = 3.5 \cdot 10^{-6}$). Furthermore, we found that the number of clicks decreased monotonically with equifinality (Spearman's $\rho = -0.98, p = 0.0004$) and so did the number of means selected by the participants (Spearman's $\rho = -1, p < 0.0001$). We also found that the median reaction time decreased significantly with equifinality ($\rho = -0.98, p = 0.00$).

Finally, we tested whether the effect of equifinality on relative performance was mediated by the random solution likelihood using the mediation analysis methodology by Baron and Kenny (1986). We found that the effect of equifinality on solution quality was not mediated by the random solution likelihood because the random solution likelihood had no significant effect on solution quality ($F(398) = 1.79, p = 0.18$). However, the random solution likelihood did have a significant effect on solution frequency ($\beta_{RSL} = 0.20, t(398) = 5.29, p = 2.8 \cdot 10^{-8}$), and this effect fully mediated the effect of high equifinality (i.e., $d = 12$ vs. $d \leq 10$) on

solution frequency: On its own, high equifinality had a significant effect on solution frequency ($\hat{\beta} = 1.28, t(396) = 2.02, p = 0.04$). Furthermore, high equifinality also had a significant effect on the random solution likelihood ($\hat{\beta} = 5.05, t(398) = 12.89, p < 10^{-15}$). Critically, when we controlled for the effect of random solution likelihood, then the effect of high equifinality on solution frequency was no longer statistically significant ($\hat{\beta} = 0.58, t(395) = 0.87, p = 0.38$). This suggests that the increase in the solution frequency from the conditions with $d \leq 10$ to the condition with $d = 12$ was due to the concomitant increase in the proportion of optimal solutions.

The size of the optimal solution was positively correlated with the solution quality ($\rho = 0.15, p = 0.0025$) but negatively correlated with whether or not the optimal solution was found ($\rho = -0.13, p = 0.01$).

These results support our theoretical prediction that people’s ability to identify an optimal set of means for achieving their goals is a U-shaped function of equifinality (Hypothesis II).

Experiments 4 and 5 had two limitations: First, multifinality and equifinality were confounded with visual clutter in the display of the connections between means and goals. Second, the goals and means were presented in abstract terms, such that participants perhaps might not have processed them as such. To address these concerns, we ran a follow-up experiment with the following, different methodology.

4.5 Experiment 6: Multifinality and human performance in a semantic Set Cover task

4.5.1 Methods

4.5.1.1 Participants. We recruited 50 participants on Amazon Mechanical Turk.

Participants were paid \$3.25 and could earn a bonus of up to \$4. The median completion time of the experiment was 45 minutes, and the average bonus was \$0.91.

4.5.1.2 Stimuli and Procedure. To address the concern about visual clutter we removed the edges from means to goals and instead listed the goals achieved by each means next to its checkbox (see Figure 10). Namely, in our interface, when a subject selected a mean, the set of all goals that could be covered if this means were chosen was highlighted. We had a checkmark next to every goal that was already covered by previously-chosen means. To address the second concern, we added the meaningful labels introduced in Experiment 3. To see whether the concerns were justified we reran Experiment 4 with this methodology and determined whether any of the conclusions changed as a result.

The experimental design was a within-subjects design with 8 conditions where the number of goals achieved by each means was 2, 3, 4, 5, 6, 8, 10, and 12 respectively. We used the same graphs from Experiment 4 to generate the SC problem participants were asked to solve. On each trial participants were required to achieve all goals in order to submit their solution.

The consent form specified that participants must not have participated in the previous version of this experiment. Participants completed one practice trial before solving the eight problems. For each problem, participants were incentivized by a bonus of \$0.50 for solving it optimally.

4.5.2 Results

Participants perceived the problems to be challenging but were highly motivated to solve them: Participants rated the difficulty of tasks as 7.36 ± 1.48 (means $\pm$ SD) on a 9-point scale. Their self-reported motivation to find the optimal

solution was 7.36 ± 1.77 , and the average rating for the motivation to finish the tasks as quickly as possible was only 4.16 ± 2.09 .

[INSERT FIGURE 13 HERE]

First, we analyzed the effect of multifinality on solution frequency and solution quality (see Figure 13). We found that the frequency with which people found an optimal solution differed significantly between the 8 groups ($\chi^2(7) = 42.5, p < 10^{-6}$). Next, we investigated whether this relationship is U-shaped. We found that, while d has a significant negative effect on solution frequency ($\beta_d = -0.92, t(397) = -4.86, p = 1.2 \cdot 10^{-6}$), d^2 had a significant positive effect on it ($\beta_{d^2} = 0.06, t(397) = 4.85, p = 1.2 \cdot 10^{-6}$). The additional variance explained by adding d^2 was statistically significant ($F(1, 397) = 27.00, p = 3.3 \cdot 10^{-7}$). Next, we tested whether the minimum point lies inside the range of degrees we investigated according to the method proposed by Lind and Mehlum (2010). The model's estimate for the minimum point is $d_{\min} = \frac{-\hat{\beta}_d}{2 \cdot \hat{\beta}_{d^2}}$. Therefore, we tested the alternative hypothesis that $\beta_d + 2 \cdot x_L \cdot \beta_{d^2} < 0$ and $\beta_d + 2 \cdot x_H \cdot \beta_{d^2} > 0$ against the null hypothesis which says that at least one of the two statements is false. Based on the results for the solution quality we tested whether the minimum lies between 6 and 9. In support of this hypothesis, we found that both components of the null hypothesis can be rejected ($F_1(1, 397) = 21.66, p < 10^{-5}$, $F_2(1, 397) = 14.63, p = 0.0002$) for $x_L = 4$ and $x_{\text{high}} = 9$. We therefore conclude that multifinality has a U-shaped effect on solution frequency.

Likewise, the average solution quality differed significantly between the five conditions ($F(7, 392) = 10.87, p < 10^{-11}$). To test our hypothesis, that average

solution quality is a U-shaped function of the number of goals per mean, we performed a fitted a general linear model with the regressors d and d^2 . We found that, while d had a significant negative effect on solution quality ($\beta_d = -0.06$, $t(397) = -6.8$, $p < 10^{-10}$), d^2 had a significant positive effect on the average solution quality ($\beta_{d^2} = 0.004$, $t(397) = 5.8$, $p < 10^{-7}$) and the additional variance explained by d^2 above and beyond the variance explained by d was highly significant ($F(1,397) = 33.35$, $p = 1.6 \cdot 10^{-8}$). To test whether the non-linear effect of d was U-shaped, we estimated the location of the minimum point (Lind & Mehlum, 2010; Bebu, Seillier-Moiseiwitsch, & Mathew, 2009). According to our data we are 99% confident that the minimum point lies between 7.67 and 9.78 ($\hat{d}_{min} = 8.37$, 99% CI: [7.6671, 9.7757]). We can therefore conclude that the average solution quality is a U-shaped function of multifinality. Furthermore, we found that the number of clicks decreased monotonically with multifinality (Spearman's $\rho = -0.95$, $p = 0.0011$) and so did the number of means selected by the participants (Spearman's $\rho = -1$, $p < 0.0001$).

The size of the minimal solution was positively correlated with people's relative performance ($\rho = 0.40$, $p < 0.0001$) and with whether or not they found the optimal solution ($\rho = 0.13$, $p = 0.01$).

The results of this experiment replicate the findings from Experiment 4 and 5. This suggests that the observed effects of multifinality and equifinality observed in Experiments 4 and 5 were not due to the confounding effect of visual clutter or the lack of meaningful labels.

4.6 Discussion of equifinality and multifinality experiments

In three experiments, we systematically manipulated the average degree of a goal system. In Experiments 4 and 6, we examined how increasing the degree of means affects human performance. Experiment 5 examined how changing the degree of goals influences performance. In all three experiments, increasing the degree resulted in a U-shaped relationship: Initially, increasing the degree decreased performance. As the degree (either of goal or means vertices) approached the number of means (or number of goals) the task of finding an optimal set of means became easier and the performance of people gradually improved. As before, a similar effect emerged in two different visual displays, both when goal vertices were given or absent a semantic content.

Previous research (Zhang, Fishbach, & Kruglanski, 2007) has examined how the degree of a *single* means influences its perceived instrumentality to goal attainment. It was shown that, as a means is connected to more goals, it is perceived as less instrumental to goal achievement and is less likely to be chosen. Since we kept the degree of all means-vertices identical in each condition in Experiments 4 and 6, it appears unlikely that the reduced performance in the Set Cover task was a result of decreased motivation. While the degree of the means varied within the conditions of Experiment 5, there was no statistically significant correlation between the number of goals connected to a means and the frequency with which it was selected in any of these conditions. This suggests that the results of Experiment 5 cannot be explained by postulating that participants were biased towards selecting means connected to fewer goals.

There is also evidence that increasing the number of means connected to a goal may decrease the motivation to pursue this goal (Huang & Zhang, 2013) once

substantial progress has been made towards goal attainment. In contrast, the presence of multiple means in the beginning of pursuing the goal enhances motivation and makes the goal appear more attainable. In our studies, we increased the average degree of entire goal system, and people were asked to cover all goals (as opposed to choosing between goals with differing number of means). Furthermore, the number of means was presented to them from the very beginning. Finally, it seems highly unlikely that the monotonic decrease of motivation with the degree of means and goals would be sufficient to explain the non-monotonic, U-shaped effects of equi- and multi-finality on human performance. To summarize, we find it unlikely that our results follow from reduced motivation in goal-systems that are denser.

From a normative perspective, increasing both the multifinality and equifinality of a goal system results in greater efficiency as it often entails that one can cover all goals with a fewer means. Interestingly, while our participants were able to achieve their goals with fewer means as the average degree of the goal system increased (all $\rho \geq 0.99$ and all $p \geq 0.02$), it became harder for them to find the optimal or near optimal solution. Namely, the size of the optimal solution was positively correlated with our two performance metrics. These results indicate that boundedly rational agents which are limited to efficient computation may fail to take advantage of a denser goal system especially if its success in goal pursuit hinges of finding optimal or near optimal solutions.

In our analysis here, we do not consider additional predictors such as our tree-likeness measures from Experiments 1-3. This was done primarily due to practical constraints as the number of graphs used in Experiments 4-6 was large and computing many of the tree-likeness parameters is NP-hard and consumes prohibitive amount of

time. It is possible that the effect of the degree of performance was (partly) mediated by tree-likeness as increasing the degree can increase treewidth and expansion. We leave it for future work to examine possible connections between our tree-likeness measures and the average degree.

5. General discussion

In this article, we investigated why people sometimes fail to pursue their goals effectively even when they are highly motivated and highly qualified to achieve them. We hypothesized that this happens because the complex structure of a person's goal system can render it prohibitively difficult for them to choose the right set of means to optimally pursue all of their goals simultaneously. We tested this hypothesis in a series of behavioral experiments and found that theoretically-motivated structural properties of goal systems are highly predictive of human performance. Our findings were consistent across two different means-selection tasks (Maximum-Coverage and Set Cover), two families of graph-theoretic metrics (similarity to trees and average degree), and two different visual displays. All our findings were in line with the complexity-theoretic predictions we derived from the literature in theoretical computer science. This suggests that computational complexity theory and insights from algorithm design can be useful theoretical tools for understanding what makes tasks cognitively difficult.

5.1 Implications for motivation and goal pursuit

The empirical results presented here indicate that the combinatorial optimization problems that people need to solve to achieve their goals can be a critical bottleneck to success regardless of motivation. Psychological theories seeking to

explain failures of self-control and motivation should therefore acknowledge the computational problems that people need to solve to figure out what they have to do to pursue their many goals as effectively as possible. Thus, predictions of how successful somebody will be should not solely be based on their motivation but shall also take into account the structure of their goal system, as even highly motivated people may perform poorly when the computational burden involved in goal pursuit is high. As far as we know, this is the first article to investigate the bounded rationality of human goal pursuit from a computational perspective. Elucidating which heuristics people use to cope with the computational challenges entailed by pursuing multiple goals in parallel is an interesting direction for future research.

When people pursue goals “in the wild”, their goals are likely to be represented differently than in our experiments. At a given moment, people may not be aware of all of their goals, all the means that are instrumental to their goals, and all connections between the desired goals and the available means. It is also likely that people’s limited working memory capacity (Miller, 1956) likely makes it impossible for them to maintain the entirety of their goal system in working memory when deciding what to do. Hence, real-life representations of goals are likely to make means-selection problems even more challenging to solve. The additional difficulty of maintaining goal systems in working memory along with the uncertainties of connection of goals and means strengthens our conclusion that a major source of difficulties in goal pursuit hinges in the computational hardness of the problems involved.

A possible direction for future investigations would be to measure people’s goal systems with an extension of personal projects analysis (Little & Gee, 2007) as well as extract graph-theoretic metrics to characterize their structure and measure the

computational complexity of pursuing goals in real life. Building on this methodology, future work could examine whether goal systems of people that are more successful in achieving their goals differ structurally from those of people that finding harder to accomplish their aims. This approach would also enable more ecologically valid experiments that study means-selection in simulations of real-life goal pursuits. Since our method allows us to predict goal achievement from the structure of person's goal system, it could also be used to detect overly complicated goal systems that set people up for failure and help them restructure their goals to make it easier to achieve them. Our findings suggest that this could be achieved by reducing the number of goals, reducing the number of goals that are considered to evaluate each means or the number of means that are considered relevant to each goal, or reducing the inter-dependencies between different goals. Alternatively, computerized tools could be used to help people deal with their goal systems (although, computers are, to some extent, constrained by intractability considerations as well).

5.2 Implications for cognitive science

As noted, NP-hard problems are abundant in various domains such as perception, decision making, language and communication, categorization and problem solving. Our results suggest that computational intractability can constrain the algorithms people use to solve such challenging problems. Furthermore, intractability results can be used to reason about cognitive and behavioral outcomes of bounded rational agents who face difficult problems. In particular, structural properties can be used to reason about whether instances faced by the cognitive system are likely to be hard or easy to solve. This research program can be seen as a

version of rational analysis (Anderson, 1990) that characterizes the structure of the problems people face in terms of properties that can be exploited by specialized algorithms. These insights could then be used to derive which algorithms the mind should use to solve those problems as efficiently as possible and evaluate those predictions in dedicated experiments (Griffiths, Lieder, & Goodman, 2015). Such future experiments could yield new theoretical insights about how the mind works and shed some light on how computationally constrained agents solve difficult problems. These insights in turn could be useful in designing more effective AI systems.

5.3 Limitations

A limitation of the experiments reported here is that they consider a restricted form of goal systems. To allow clean and simple examination of our hypotheses we focused on unweighted networks where all goals have the same value and all means have the same cost. However, in real life, some goals are much more valuable than others and some means are much more expensive than others. Future experiments should investigate the predictive power of computational complexity theory in more realistic settings where the goals differ in their values and the means differ in their costs¹¹. As in goal-systems theory we are mostly concerned with shallow goal systems, it is quite likely that people encounter goal systems of depth greater than 2. Examining empirically how people cope with such deeper systems is of potential

¹¹ In a follow-up experiment, we found that tasks where the goals differ in value are significantly more difficult than tasks where all goals have the same value. Concretely, when people were tasked to select five means so as to maximize the sum of the achieved goals' values, the frequency of the optimal solution dropped to 7%. This suggests that, as goal management tasks become more realistic, their computational difficulty becomes even more limiting. Due to low average performance in this challenging task, the effect of treewidth on solution frequency was no longer detectable ($\rho = -0.18, p = 0.46$). This suggests that treewidth is not the only factor determining computational difficulty.

interest. Future research could also consider problems where multiple means have to be invested to achieve a goal and some means contribute more towards achieving a goal than others (namely goal systems and goal systems with weighted edges). In particular, one would expect that in certain circumstances some means would inhibit some goals (Kruglanski, Chernikova, Babush, Dugas & Schumpe, 2015). Hence, it could be of interest to examine how people perform on goal systems with negative weights. Finally, goal pursuit often entails sequential decisions and actions over time. Looking into such sequential tasks and how people deal with the computational challenges they entail is another avenue for future research.

5.4 Conclusion

Striving for goals is essential to the human experience, and progress towards one's goals is critical to attain satisfaction and well-being. The work presented here sheds new light on this fundamental human problem and the factors which determine whether it is easy or difficult. Contrary to previous accounts, our theory suggests that effective goal pursuit is not just a matter of motivation but also entails very challenging combinatorial optimization problems. Our results suggest that whether a person will be able to achieve their goals partly depends on the structure of their goal system, as this structure determines how cognitively demanding it is to find the right course of action. This line of work will advance our basic understanding of human performance, and it might also lead to novel interventions and software tools for helping people restructure their goals so that they can be pursued more effectively.

Acknowledgements

We thank Rina Dechter for advise on treewidth computation and Kevin Leyton-Brown, Nira Liberman, and Rika Antonova for helpful comments and discussions. This research was supported by Grant ONR MURI N00014-13-1– 0341 from the Office of Naval Research to Thomas L. Griffiths.

References

- Ainslie, G., & Haslam, N. (1992). Self-control. In G. Loewenstein and J. Elster (Eds.). *Choice over time* (pp. 177—209). New York: The Sage Foundation.
- Alber, J., & Niedermeier, R. (2002, April). Improved tree decomposition based algorithms for domination-like problems. In *Latin American Symposium on Theoretical Informatics*, (pp. 613-627). Berlin: Springer-Verlag.
- Alon, N., & Milman, V. D. (1985). λ_1 , isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory, Series B*, 38(1), 73-88.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Ariely, D., & Wertenbroch, K. (2002). Procrastination, deadlines, and performance: using precommitment to regulate one's behavior. *Psychological Science*, 13, 219-224.
- Austin, J. T., & Vancouver, J. B. (1996). Goal constructs in psychology: Structure, process, and content. *Psychological Bulletin*, 120, 338–375.
- Baker, B. S. (1994). Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM*, 41(1), 153-180.
- Bandura, A. (1997). *Self-efficacy: The exercise of control*. New York, NY: W. H. Freeman/Times Books/Henry Holt & Co.
- Baron, R. M., & Kenny, D. A. (1986). The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of personality and social psychology*, 51(6), 1173.

Baumeister, R. F., & Heatherton, T. F. (1996). Self-regulation failure: An overview.

Psychological Inquiry, 7, 1–15.

Bebu, I., Seillier-Moiseiwitsch, F., & Mathew, T. (2009). Generalized confidence intervals for ratios of regression coefficients with applications to bioassays.

Biometrical Journal, 51(6), 1047-1058.

Bodlaender, H. L. (1988, July). Dynamic programming on graphs with bounded treewidth. In *International Colloquium on Automata, Languages, and Programming* (pp. 105-118). Berlin: Springer-Verlag.

Bossaerts, P., & Murawski, C. (2017). Computational complexity and human decision-making. *Trends in Cognitive Sciences*, 21(12), 917-929.

Cantor, N., & Blanton, H. (1996). Effortful pursuit of personal goals in daily life. In J. A. Bargh & P.M. Gollwitzer (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 338–359). New York, NY: Guilford.

Carruthers, S., Masson, M. E., & Stege, U. (2012). Human performance on hard non-Euclidean graph problems: Vertex cover. *The Journal of Problem Solving*, 5(1), 34-55.

Carver, C. S., & Scheier, M. F. (1998). *On the self-regulation of behavior*. New York, NY: Cambridge University Press.

Clementi, A. E., & Trevisan, L. (1999). Improved non-approximability results for minimum vertex cover with density constraints. *Theoretical Computer Science*, 225(1), 113-128.

- Cormen, T. H., Leiserson, C. E., Rivest, R. L., & Stein, C. (2001). *Introduction to algorithms* (2nd ed.). Cambridge, MA: MIT Press.
- DeShon, R. P., & Gillespie, J. Z. (2005). A motivated action theory account of goal oriented behavior. *Journal of Applied Psychology*, 90, 1096–1127.
- Dinur, I., Guruswami, V., Khot, S., & Regev, O. (2005). A new multilayered PCP and the hardness of hypergraph vertex cover. *SIAM Journal of Computing*, 34, 1129-1146.
- Dworkin, L., & Kearns, M. (2015). From "in" to "over": Behavioral experiments on whole-network computation. In *Third AAAI Conference on Human Computation and Crowdsourcing* (pp. 52 - 61). Palo Alto, CA: AAAI Press.
- Enemark, D. P., McCubbins, M. D., Paturi, R., & Weller, N. (2011, June). Does more connectivity help groups to solve social problems?. In *Proceedings of the 12th ACM Conference on Electronic Commerce* (pp. 21-26). New York, NY: ACM Press.
- Feige, U. (1998). A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4), 634-652.
- Fishbach, A., Shah, J. Y., & Kruglanski, A. W. (2004). Emotional transfer in goal systems. *Journal of Experimental Social Psychology*, 40(6), 723-738.
- Fishbach, A., Zhang, Y., & Koo, M. (2009). The dynamics of self regulation. *European Review of Social Psychology*, 20, 315–344.
- Garey, M. R., & Johnson, D.S. (1979). *Computers and intractability: A guide to the theory of NP-completeness*. New York: Freeman.

Gigerenzer, G. (2008). *Rationality for mortals*. New York: Oxford University Press.

Gogate, V., & Dechter, R. (2004). A complete anytime algorithm for treewidth. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence* (pp. 201-208). Arlington, VA: AUAI Press.

Goldreich, O. (2008). *Computational complexity: A conceptual perspective*. New York, NY: Cambridge University Press.

Goldreich, O. (2010). *P, NP, and NP-completeness: The basics of complexity theory*. New York, NY: Cambridge University Press.

Gollwitzer, P. M. (1999). Implementation intentions: strong effects of simple plans. *American Psychologist*, 54, 493–503.

Green, L., Fristoe, N., & Myerson, J. (1994). Temporal discounting and preference reversals in choice between delayed outcomes. *Psychonomic Bulletin & Review*, 1(3), 383-389.

Griffiths, T.L., Lieder, F., & Goodman, N.D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7(2), 217-229.

Harel, D., & Feldman, Y. (2004). *Algorithmics: The spirit of computing* (3rd ed.). Berlin: Springer-Verlag.

Huang, S. C., & Zhang, Y. (2013). All roads lead to Rome: The impact of multiple attainment means on motivation. *Journal of Personality and Social Psychology*, 104(2), 236-248.

- Hutter, F., Xu, L., Hoos, H. H., & Leyton-Brown, K. (2014). Algorithm runtime prediction: Methods & evaluation. *Artificial Intelligence*, 206, 79-111.
- Iwata, Y. (2016). Linear-time kernelization for feedback vertex set. *arXiv preprint arXiv:1608.01463*.
- Johnson, R. E., Chang, C-H. D., & Lord, R. G. (2006). Moving from cognition to behavior: What the research says. *Psychological Bulletin*, 132, 381-415.
- Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice-Hall.
- Kask, K., Gelfand, A., Otten, L., & Dechter, R. (2011). Pushing the power of stochastic greedy ordering schemes for inference in graphical models. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence* (pp. 54-60). Palo Alto, CA: AAAI Press.
- Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220(4598), 671-680.
- Klein, P., Plotkin, S. A., & Rao, S. (1993). Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing* (pp. 682-690). New York, NY: ACM Press.
- Kleinberg, J., & Tardos, E. (2005). *Algorithm design*. London, UK: Pearson Education.
- Kloks, T. (1994). *Treewidth: Computations and approximations* (Vol. 842). New York: Springer Science & Business Media.

- Köpetz, C., Faber, T., Fishbach, A., & Kruglanski, A. W. (2011). The multifinality constraints effect: How goal multiplicity narrows the means set to a focal end. *Journal of Personality and Social Psychology*, 100(5), 810.
- Kruglanski, A. W., Bélanger, J. J., Chen, X., Köpetz, C., Pierro, A., & Mannetti, L. (2012). The energetics of motivated cognition: A force-field analysis. *Psychological review*, 119(1), 1.
- Kruglanski, A. W., Chernikova, M., Babush, M., Dugas, M., & Schumpe, B. (2015). The architecture of goal systems: Multifinality, equifinality, and counterfinality in means-end relations. *Advances in Motivation Science*, 2, 69-98.
- Kruglanski, A. W., Shah, Y. J., Fishbach, A., Friedman, R., Chun, W. Y., & Sleeth-Keppler, D. (2002). A theory of goal-systems. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (pp. 331–378). San Diego, CA: Academic Press.
- Kruglanski, A. W., & Köpetz, C. (2009). What is so special (and non-special) about goals? A view from the cognitive perspective. In G. B. Moskowitz & H. Grant (Eds.), *The psychology of goals* (pp. 27-55). New York: Guilford Press.
- Kruglanski, A.W., & Köpetz, C. (2010). The self control challenge in human action: Unpacking the dilemma. In R. Hassin, K. Ochsner, & Y. Trope (Eds.), *Self-control in society, mind, and brain* (pp. 297-312). Oxford, UK: Oxford University Press.
- Leyton-Brown, K., Nudelman, E., & Shoham, Y. (2009). Empirical hardness models: Methodology and a case study on combinatorial auctions. *Journal of the ACM*, 56(4), 1-52.

- Leyton-Brown, K., Hoos, H. H., Hutter, F., & Xu, L. (2014). Understanding the empirical hardness of NP-complete problems. *Communications of the ACM*, 57(5), 98-107.
- Liberman, N., & Dar, R. (2009). Normal and pathological consequences of encountering difficulties in monitoring progress toward goals. In G. B. Moskowitz and H. Grant (Eds.), *The psychology of goals* (pp. 277-303). New York: Guilford.
- Lind, J. T., & Mehlum, H. (2010). With or without U? the Appropriate Test for a U-Shaped Relationship. *Oxford bulletin of economics and statistics*, 72(1), 109-118.
- Little, B. R., & Gee, T. (2007). The methodology of personal projects analysis: Four modules and a funnel. In B. R. Little, K. Salmela-Aro, & S. D. Philipps (Eds.), *Personal project pursuit: Goals, action, and human flourishing*. New York: Psychology Press.
- Louro, M. S., Pieters, R., & Zeelenberg, M. (2007). Dynamics of multiple goal pursuit. *Journal of Personality and Social Psychology*, 93, 174-193.
- Lovász, L. (1975). On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13(4), 383-390.
- MacGregor, J. N., & Ormerod, T. (1996). Human performance on the traveling salesman problem. *Attention, perception, & psychophysics*, 58(4), 527-539.

- MacGregor, J. N., Ormerod, T. C., & Chronicle, E. P. (2000). A model of human performance on the travelling salesperson problem. *Memory and Cognition*, 28, 1183–1190.
- Marx, D. (2007, October). Can you beat treewidth?. In *48th Annual IEEE Symposium on Foundations of Computer Science*. (pp. 169-179) Washington, DC: IEEE Computer Society.
- McGregor, I., & Little, B. R. (1998). Personal projects, happiness, and meaning: on doing well and being yourself. *Journal of Personality and Social Psychology*, 74(2), 494.
- Miller, G. A. (1956). The magic number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 91-97.
- Miller, G.A., Galanter, E., & Pribram, K.H. (1960). *Plans and the structure of behavior*, New York, NY: Holt.
- Mischel, W. (1974). Processes in delay of gratification. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 7, pp. 249-292). San Diego: Academic Press.
- Mischel, W., Shoda, Y., & Peake, P. K. (1988). The nature of adolescent competencies predicted by preschool delay of gratification. *Journal of Personality and Social Psychology*, 54, 687-696.
- Muraven, M., & Baumeister, R.F. (2000). Self-regulation and depletion of limited resources: Does self-control resemble a muscle? *Psychological Bulletin*, 126, 247–259.

- Myerson, J., & Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the Experimental Analysis of Behavior*, 64(3), 263-276.
- Newell, A., Shaw, J.C., & Simon, H.A. (1958). Elements of a theory of human problem solving. *Psychological Review*, 65(3), 151-166.
- Papadimitriou, C. H (1994). *Computational complexity*. Chichester, UK: Addison-Wesley.
- Perarnau, G., & Serra, O. (2014). On the tree-depth of random graphs. *Discrete Applied Mathematics*, 168, 119-126.
- Robertson, N., & Seymour, P. D. (1986). Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3), 309-322.
- Saket, R., & Sviridenko, M. (2012). New and improved bounds for the minimum set cover problem. In *Approximation, randomization, and combinatorial optimization. Algorithms and techniques* (pp. 288-300). Berlin: Springer.
- Sanborn, A., Griffiths, T., & Navarro, D. J. (2010). Rational approximation to rational models: Alternative algorithms for category learning. *Psychological Review*, 117(4), 1144–1167.
- Shah, J.Y., & Kruglanski, A.W. (2000). Aspects of goal networks: Implications for self-regulation. In M. Boekaerts and P.R. Pintrich (Eds.). *Handbook of self-regulation*. (pp. 85-110). San Diego: Academic Press.

- Shah, J.Y., Friedman, R., & Kruglanski, A.W. (2002). Forgetting all else: On the antecedents and consequences of goal shielding. *Journal of Personality and Social Psychology*, 83, 1261–1280.
- Shah, J., Hall, D., & Leander, N. P. (2009). Moments of motivation: Margins of opportunity in managing the efficacy, need and transitions of striving. In G. Moskowitz, & H. Grant (Eds.), *The psychology of goals* (pp. 234–254). New York: The Guilford Press.
- Shah, J. Y., & Kruglanski, A. W. (2003). When opportunity knocks: Bottom-up priming of goals by means and its effects on self-regulation. *Journal of Personality and Social Psychology*, 84, 1109-1122.
- Sheldon, K. M., & Elliot, A. J. (1999). Goal striving, need satisfaction, and longitudinal well-being: The self-concordance model. *Journal of Personality and Social Psychology*, 76(3), 482.
- Stajkovic, A. D., & Luthans, F. (1998). Self-efficacy and work-related performance: A meta-analysis. *Psychological bulletin*, 124(2), 240.
- Steel, P. (2007). The nature of procrastination: A meta-analytic and theoretical review of quintessential self-regulatory failure. *Psychological Bulletin*, 133, 65-94.
- Thagard, P., & Verbeurgt, K. (1998). Coherence as constraint satisfaction. *Cognitive Science*, 22, 1–24.
- Trevisan, L. (2001, July). Non-approximability results for optimization problems on bounded degree instances. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing* (pp. 453-461), New York: ACM Press.

- van Rooij, I. (2008). The tractable cognition thesis. *Cognitive Science*, 32, 939–984.
- van Rooij, I., Stege, U., & Kadlec, H. (2005). Sources of complexity in subset choice. *Journal of Mathematical Psychology*, 49(2), 160-187.
- van Rooij, I. & Wareham, T. (2012). Intractability and approximation of optimization theories of cognition. *Journal of Mathematical Psychology*, 56, 232-247.
- van Rooij, I., Wright, C., & Wareham, H.T. (2012). Intractability and the use of heuristics in psychological explanations. *Synthese*, 187, 471-487.
- Vazirani, V. V. (2003). *Approximation algorithms*. Atlanta, GA: Springer.
- Vroom, V. H. (1964). *Work and motivation*. New York, NY: Wiley.
- Wahlström, M. (2014). Half-integrality, LP-branching and FPT algorithms. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms* (pp. 1762-1781). New York: ACM Press.
- Wason, P.C. (1966). Reasoning. In Foss, B. M. (ed.). *New horizons in psychology*. Harmondsworth: Penguin
- Webb, T., & Sheeran, P. (2006). Does changing behavioral intentions engender behavior change? A meta analysis of experimental evidence. *Psychological Bulletin*, 132, 249–268.
- Zhang, Y., Fishbach, A., & Kruglanski, A. W. (2007). The dilution model: How additional goals undermine the perceived instrumentality of a shared path. *Journal of Personality and Social Psychology*, 92, 389-401.

Figure 1. Example of goal system: Means appear at the bottom marked by $m_1, \dots, m_4$. Goals appear on top marked by $g_1, \dots, g_5$. A goal can be attained if at least one means connected to it is chosen. The set $\{m_1, m_3\}$ is a minimal set of means covering all goals.

Figure 2. Treewidth. Some examples of graphs and their treewidths. A. Trees (upper left) have treewidth 1. B. Rectangular grids of dimension d have treewidth d , hence this graph has treewidth 3. C. Complete graphs of order n where every two pairs of nodes are connected have treewidth $n-1$. Hence the treewidth of this graph is 3.

Figure 3. A. A sample graph G . B. Two possible tree decompositions of G , the first with treewidth 2, the second with treewidth 3. The definition of tree decomposition implies that if we remove all vertices residing within a bag of the tree T (where T is a tree decomposition of G), then no two vertices of G will have an edge connecting them. Just as in the case of a tree, removing all vertices in a bag breaks the graph G into disjoint connected components.

Figure 4. The edge expansion of a graph $G=(V,E)$ is the minimal edge expansion over all of its vertex subsets whose size is at most $|V|/2$. Bounded degree trees have small edge expansion. In this figure, we consider the circled subset S of 4 vertices in a 4-regular graph. Edges that connect vertices in S to vertices not belonging to S are highlighted in red. The edge expansion of S is $8/4=2$, since there are 8 edges “going out” of S , and we normalize by the size of S .

Figure 5. Interface and instructions for Experiment 1.

a) Solution Frequency

b) Solution Quality

Figure 6. Human performance in the set cover task of Experiment 1 as a function of treewidth.

Figure 7. Human performance in the Maximum-Coverage task of Experiment 2 as a function of treewidth.

Well done! Now here is the real task: Please help John achieve **ALL** of his goals with **as few activities as possible**. To do so please select those and only those activities that you would advise him to do. A green checkmark on a goal means that at least one of the chosen activities will accomplish it.

If you discover how to achieve **ALL** of the goals with **as few activities as possible**, then we will pay you a bonus of \$0.5.

Activities	Serves Goals	Goal
☐ A	12, 18, 21, 23	☐ 1. Volunteer/give to charity
☐ B	2, 3, 9, 24	☐ 2. Improve mental skills (e.g., concentration)
☐ C	3, 8, 12, 21	☐ 3. Start a hobby
☐ D	8, 9, 15, 24	☐ 4. Learn about art, music, and culture
☐ E	3, 5, 12, 18	☐ 5. Get a better job
☐ F	1, 6, 9, 15	☐ 6. Get in shape
☐ G	11, 17, 20, 22	☐ 7. Start saving money

Figure 8. Interface and instructions for Experiment 3.

a) Solution Frequency

b) Solution Quality

Figure 9. Human performance in the semantic Set Cover task of Experiment 3 as a function of treewidth.

Figure 10. Three examples of goal systems. In each graph, the top row of vertices corresponds to the means while the bottom row of vertices corresponds to the goals. The three graphs differ in the number of goals connected to each means (i.e., the degree, d). The shaded vertices are the optimal solutions to the corresponding Set Cover problem.

(a)

(b)

Figure 11. Effect of multifinality on (a) the solution frequency and (b) the solution quality in the Set Cover task in Experiment 4.

Figure 12. Effect of equifinality on participants' solution quality (a) and solution frequency (b) in the Set Cover task in Experiment 5.

(a)

(b)

Figure 13. Effect of multifinality on solution frequency (a) and solution quality (b) in Experiment 6.

Appendix A. Graph theoretic definitions

In this section we provide some graph-theoretic definitions. We refer the interested reader to good sources for graph theory, such as Diestel (2000). An *undirected* graph G is composed of a finite set of vertices, V and a set of edges, E which is a subset of the family of all subsets of V which are of size exactly 2. Namely, each edge is an *unordered* pair $\{u, v\}$ where both u, v are in V . We write $G = (V, E)$ to signify a graph G that consists of a vertex set V and edge set E . We say that a vertex y is an *neighbor* of x if $\{x, y\} \in E$. Alternatively, we say that y is adjacent to x . Note that x is adjacent (a neighbor) to y if and only if y is adjacent to x . The *degree* of x is defined as the number of neighbors of x . A *simple path* is a set of *distinct* vertices $v_1, v_2, \dots, v_k$ such that for every $1 \leq i < k$, v_i is a neighbor v_{i+1} . The *length* of the path is the number of vertices in the path minus one (that is, $k - 1$). A graph is *connected* if every two vertices are connected by a path of finite length. A *simple cycle* is a set of distinct vertices $v_1, v_2, \dots, v_k v_1$ such that, for every i , v_i is adjacent to v_{i+1} and v_1 is adjacent to v_k . A connected graph with no cycles is called a *forest*. An *independent set* is a subset I of vertices that contains no edges. $G = (V, E)$ is *bipartite* if the vertex set of V is the union of two disjoint independent sets. Given a list of vertices, $v_1, v_2, \dots, v_r$ the degree sequence of these vertices is the list of the degrees of $v_1, v_2, \dots, v_r$. The average degree of a graph is the sum of the degrees normalized by the number of vertices. Finally, we comment that the sum of degrees (hence the average degree) is related to the number of edges in a graph: It is known that equals $2|E|$ where $|E|$ is the number of edges of the graph.

A *tree decomposition* of a graph $G=(V,E)$ is a tree $T=(U,F)$ and a function $\chi:U \rightarrow 2^V$ (namely, for every $u \in U$, $\chi(u)$ is subset of vertices from V). We term the subsets of vertices of V associated with the vertices of T , bags. Tree decompositions are defined by two properties: (i) every edge $(u,v) \in E$ belongs to some bag $\chi(t), t \in U$ and (ii) for every vertex $v \in V$ the bags containing v form a connected subtree of T . The *width* of the tree decomposition is one less than the cardinality of the largest bag (e.g., $\max_{t \in T} |\chi(t)| - 1$) and the treewidth of G , denoted by $\text{tw}(G)$ is the smallest width for which G has a tree decomposition. We shall omit G from this notation when clear from the context. See Figure 2 and 3 for examples and further explanations of treewidth.

For the computation of the treewidth of our graphs, we used the QUICKBB library (<http://graphmod.ics.uci.edu/group/quickbb>) which is described in Gogate and Dechter (2004). Using a cluster of machines, each containing 96GB of RAM, the treewidth computation of most graphs completed within 40 hours.

Given a graph G and a nonempty (strict) subset S of V , a *cut* $(S, V \setminus S)$ is a partition of the vertices of G to two nonempty sets S and $V \setminus S$. An edge of G *crosses* the cut if exactly one of its endpoints belong to S . We denote by $\partial(S)$ the set of edges crossing the cut $(S, V \setminus S)$ and define $N(S)$ to be the set of all vertices in $V \setminus S$ having a

neighbor in S . The vertex-expansion of G is defined as $\min_{\emptyset \neq S \subseteq V, |S| \leq |V|/2} \frac{|N(S)|}{|S|}$. The

edge-expansion of G is $\min_{\emptyset \neq S \subseteq V, |S| \leq |V|/2} \frac{|\partial(S)|}{|S|}$.

We computed both vertex and edge expansion by formulating these quantities as integer linear programs (ILPs) and used IBM's CPLEX to solve these programs.

Given an undirected graph $G = (V, E)$ over n vertices it can be represented by a symmetric matrix A_G of dimension n with 0-1 entries which is called the adjacency matrix of G . We have that $A_G(i, j) = 1$ if the i th vertex is connected to the j th vertex and 0 otherwise. The adjacency matrices of the graphs we consider are symmetric, hence have n real eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. For bipartite d -regular graphs, it is known that $\lambda_1 = d$ and $\lambda_n = -d$. The classical (discrete) Cheeger's inequality (Alon & Milman, 1985) implies that λ_2 and the edge expansion are related: The larger $d - \lambda_2$ is, the larger the edge expansion of the graph. The spectral expansion of G equals $d - \lambda_2$.

Appendix B. Problem graphs

In the following figures, shaded nodes constitute the optimal minimum cover sets

TW 13

TW 6

TW 12 *TW 11* *TW 12* *TW 11* *TW 10* *TW 8*

TW 5

TW 9

TW 7

TW 6

TW 10

TW 6 *TW 12* *TW 14* *TW 12* *TW 8*

TW 11

TW 9