

1 **How humans solve complex problems: The case of the knapsack problem**

2

3 *Short title:* How humans solve complex problems

4

5 Carsten Murawski<sup>1</sup> and Peter L. Bossaerts<sup>1,2,3,\*</sup>

6

7 <sup>1</sup>The University of Melbourne, Department of Finance, Parkville, Victoria 3010, Australia

8 <sup>2</sup>Florey Institute of Neuroscience and Mental Health, Parkville, Victoria 3052, Australia

9 <sup>3</sup>University of Utah, David Eccles School of Business, Salt Lake City, UT 84112, USA

10

11 *Classification:* SOCIAL SCIENCES (Economic Sciences), BIOLOGICAL SCIENCES (Psy-  
12 chological and Cognitive Sciences)

13

14 *Keywords:* Human decision-making, knapsack problem, combinatorial optimisation, compu-  
15 tational complexity, NP-complete

16 \*Corresponding author: Peter L. Bossaerts, The University of Melbourne, Department of Finance, Parkville,

17 Victoria 3010, Australia, phone: +61 3 8344 6912, email: peter.bossaerts@unimelb.edu.au.

18 **Abstract**

19 Life presents us with problems of varying complexity. Yet, complexity is not accounted for in  
20 any of the major theories of human decision-making. Here, we investigated experimentally how  
21 humans solve instances of the knapsack problem. This complex discrete optimisation problem  
22 is commonly encountered in every-day life at many levels of cognition such as allocation of  
23 attention, time management and financial decisions. It has also been related to intellectual dis-  
24 covery and innovation as well as several psychiatric disorders. We show that human ability  
25 to solve instances of the knapsack problem decreased rapidly with increasing complexity, as  
26 defined in computer science. At the same time, in violation of traditional economic principles,  
27 economic gain in an instance increased with difficulty, as revealed by the proportion of partici-  
28 pants who solved it. Behaviour exhibited signatures of solution algorithms developed for Turing  
29 machines, but seemed to be limited by working and episodic memories. Within each instance,  
30 most attempts failed, but at least one participant always found the solution quickly, yet often  
31 did not seem to realise. Our results cast doubt on one of the core assumptions of rational choice  
32 theory, question the effectiveness of the current patent system in fostering innovation—often  
33 regarded as a knapsack problem—and have major implications for diagnosis and treatment of  
34 mental disorders marked by maladaptive time and effort management.

35 **Significance statement**

36 Humans are thought to appeal to heuristics when solving difficult problems. There is no agree-  
37 ment yet on what makes a problem difficult for humans (as opposed to computers), and whether  
38 the strategies used for solving those problems are adapted to difficulty. We show that human  
39 problem solving ability decreased rapidly with computational complexity, as defined in com-  
40 puter science, provided it is adjusted suitably to accommodate biological constraints. These  
41 include limited computational capacity, working and episodic memories. We also found sub-  
42 stantial variation in problem solving ability. Schemes that promote problem solving, like patents  
43 or prizes, unfortunately fail to exploit this variation. Furthermore, choices defied traditional  
44 economic principles, because participants worked harder on, and earned more in, more difficult  
45 problems.

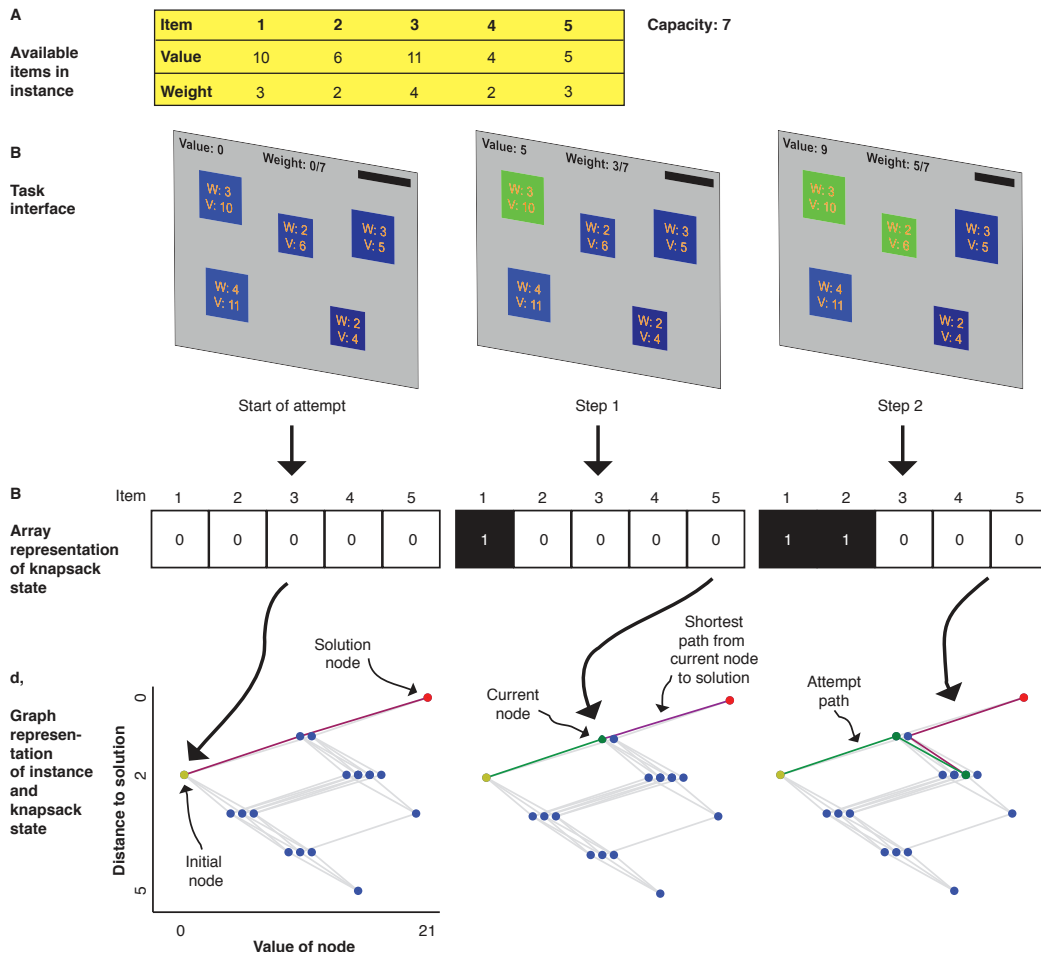
46 The knapsack problem (KP) is a combinatorial optimisation problem with the goal of finding  
47 in a set of items of given values and weights the subset of items with the highest total value,  
48 subject to a total weight constraint (1, 2) (SI Materials and Methods 1.1). It is a member of the  
49 complexity class *non-deterministic polynomial-time (NP) complete* (2–4). For those problems,  
50 there are no efficient solution algorithms, that is, algorithms whose computational time only  
51 grows as a polynomial of the size of the problem’s instances (5). While computing the value of  
52 a candidate solution—in case of the KP, the value of a set of items—is relatively easy, finding  
53 the solution—the set of items with the maximum possible value—is computationally hard.

54 The KP permeates the lives of humans (and non-human animals). It emerges at a low level  
55 of cognition, for example in the choice of stimuli to attend to (visual, auditory, tactile), referred  
56 to as optimal in-attention (6). At an intermediate level of cognition, the KP can be recognised  
57 in tasks like budgeting and time management (7). At the highest level, it occurs for example in  
58 production (cutting problems (8)), logistics (bin packing problems (9)), combinatorial auctions  
59 and in financial economics (portfolio optimisation (10)). It is also closely related to innovation.  
60 It has been argued that innovation is a combinatorial process (11–15). A recent empirical study  
61 of patent filings at the U.S. Patent Office between 1790 and 2010 showed that most patents  
62 during this period were filed for inventions that combined existing technologies in a novel way  
63 (15). Thus, innovation can be formalised as a combinatorial optimisation problem, of which the  
64 KP is a canonical example. Furthermore, the KP is related to mental disorders. For example,  
65 several of the symptoms of attention-deficit/hyperactivity disorder can be regarded as impaired  
66 ability to solve the KP (7). The KP also arises as a sub-problem in search algorithms of many  
67 integer linear programming problems (2). One reason for the ubiquity of the KP is the fact that  
68 many other NP-complete problems are reducible to the KP (16). NP-complete problems have

69 the property that if there is an efficient algorithm for one member of their class, then there  
70 exists an efficient algorithm for all members of this class as well as any problem in class NP  
71 (16), suggesting that problems of class NP-complete have one or more characteristics common  
72 to all problems in class NP. Understanding how humans approach the KP therefore promises to  
73 provide insights into a wide range of behaviours.

74 We examined the question what strategies humans use to solve complex problems like  
75 the KP and how the degree of complexity of a computational problem affects human behaviour.  
76 Twenty healthy participants attempted to solve eight instances of the 0-1 knapsack problem (13),  
77 administered on a computer (Fig. 1A, B, Materials and Methods). In each instance, participants  
78 were presented with a set of items of differing values and weights (Table S1). Participants used  
79 a mouse to add available items to and remove items from the knapsack. Total value and weight  
80 of the selected items was displayed at the top of the screen. Participants were offered two  
81 attempts per instance. The time limit per attempt was four minutes. Participants were asked to  
82 press the space bar when they believed that they had found the solution of an instance, to submit  
83 the solution. Submitted attempts were rewarded based on an economic performance measure  
84 that increased in closeness in values between the solution submitted (“attained value”) and the  
85 optimal knapsack (“solution value”).

86 The instances of the KP in our experiment differed in complexity (SI Materials and Meth-  
87 ods 1.5; Table S2). Complexity is usually defined with reference to a model of an ideal com-  
88 puter, in particular, a Turing machine (5). Measures of complexity are typically based on the  
89 number of computational steps and memory requirements of solution algorithms. It is an open  
90 question which limits the human brain imposes on a person’s ability to solve complex com-  
91 putational problems. For example, humans have very limited (working) memory and make

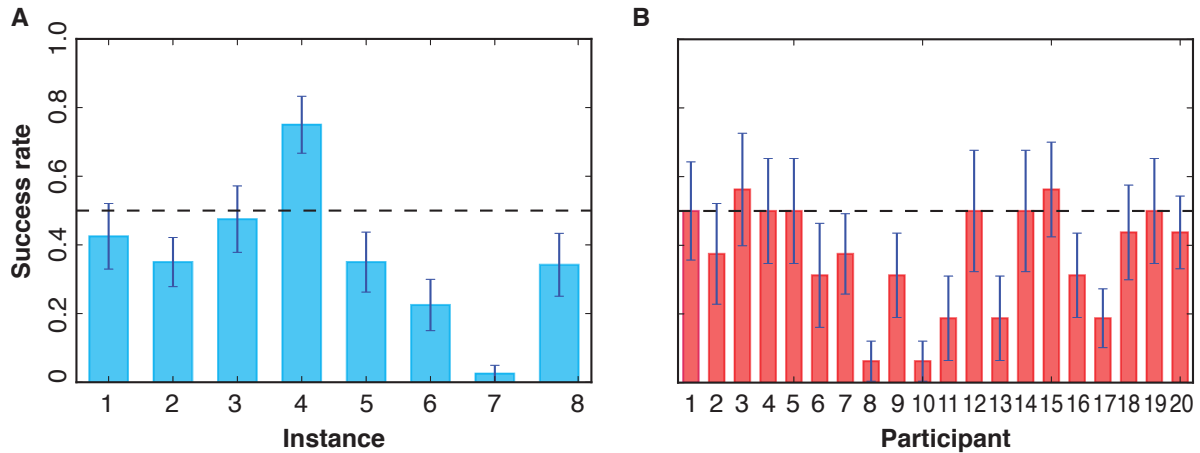


**Figure 1.** Overview of experimental paradigm. (A) Example instance of the 0-1 knapsack problem with five items. The goal is to find the subset of available items that maximises total value of the knapsack, subject to a capacity constraint. Values and weights of available items are provided in the table. The capacity of the knapsack is 7. (B) In the experiment, instances were presented on a computer interface. In the interface, each item was represented by a square. An item could be added to and removed from the knapsack by clicking on it. Selected items were displayed in green. (C) The state of the knapsack in an instance can be represented by an array of 0s and 1s with length equal to the number of items available. (D) The admissible states of the knapsack in an instance can be represented as a graph. In the graph plots, the position of a vertex on the abscissa is determined by the value of the knapsack represented by the vertex, and the position on the ordinate is determined by the distance of the vertex from the solution vertex in item space,  $d_G$  (SI Materials and Methods 1.2). The sequence of items added to and removed from the knapsack during an attempt can be represented by a path in the graph.

92 computational mistakes (17), and we would expect those features to affect their computations,  
93 at least in some domains. Our experiment was aimed at understanding in which ways and to  
94 what extent complexity of the instances in the KP affected behaviour. To this end, we tested  
95 whether various measures of complexity, proportional to the number of computational steps and  
96 memory requirements in the various instances, were related to performance in the instances.

97 **Computational performance and economic performance.** Computational performance, that  
98 is, the proportion of attempts (“score”) of an instance in which participants found the maximum  
99 value, was 37.4%. The mean time on task at submission was 172 s ( $SD = 57$  s), and 97.5%  
100 of attempts were submitted before the time limit of four minutes. While the scores indicate  
101 that our instances were difficult indeed, there was substantial variability across instances ( $min$   
102  $= 0.027$ ,  $M = 0.367$ ,  $max = 0.744$ ,  $SD = 0.19$ ; Fig. 2A), and across participants ( $min = 0.062$ ,  
103  $M = 0.374$ ,  $max = 0.562$ ,  $SD = 0.157$ ; Fig. 2B). Computational performance was significantly  
104 higher than what participants would have achieved by implementing stochastic search (‘trial  
105 and error’). Participants on average submitted only 14 different sets of items ( $SD = 7$ ), and the  
106 chance that any capacity-constrained knapsack was the optimal one was a mere 0.7%, implying  
107 an expected computational performance from random search equal to only 0.10 ( $SD = 0.05$ ).  
108 The total number of correct attempts was significantly above chance ( $P < 0.001$ , one-sided  
109 binomial test). This suggests that searches were systematic, that is, based on a non-random  
110 protocol.

111 On average, participants’ attained value equaled 97.4% of solution value. It was signif-  
112 icantly higher than the expected economic performance of an algorithm that randomly picks  
113 a capacity-constrained knapsack, which was 85.3% ( $P < 0.001$ , one-sample t-test,  $t_{(307)} =$



**Figure 2.** Variation in performance. (A) Success rates (proportion of successful attempts) for each of the eight instances administered in the task (see Table S1 for instance properties). (B) Success rates for each of the 20 participants. Blue lines indicate standard errors of means.

114 36.382). Similar to computational performance, economic performance varied more by in-  
 115 stance ( $min = 95.8\%$ ,  $M = 97.4\%$ ,  $max = 99.0\%$ ,  $SD = 1.1\%$ ) than by participant ( $min = 88.9\%$ ,  
 116  $M = 97.4\%$ ,  $max = 99.3\%$ ,  $SD = 2.4\%$ ).

117 **Performance and effort.** Decision theory predicts that economic performance is related to  
 118 the extent of effort extended. To investigate whether this was the case in our experiment, we  
 119 examined participants' search paths, that is, the sequences of item additions/deletions in an  
 120 attempt. These sequences can be expressed as paths in a graph where vertices correspond to  
 121 possible solutions, and edges indicate addition/deletion of an item (Fig. 1C, D, Fig. S1; SI  
 122 Materials and Methods 1.2). Path length then corresponds to the number of additions/deletions  
 123 (number of steps) required to travel from the initial vertex (empty knapsack) to a particular  
 124 vertex. Terminal vertices in the graph correspond to capacity-constrained knapsacks: further  
 125 additions would violate the weight constraint.

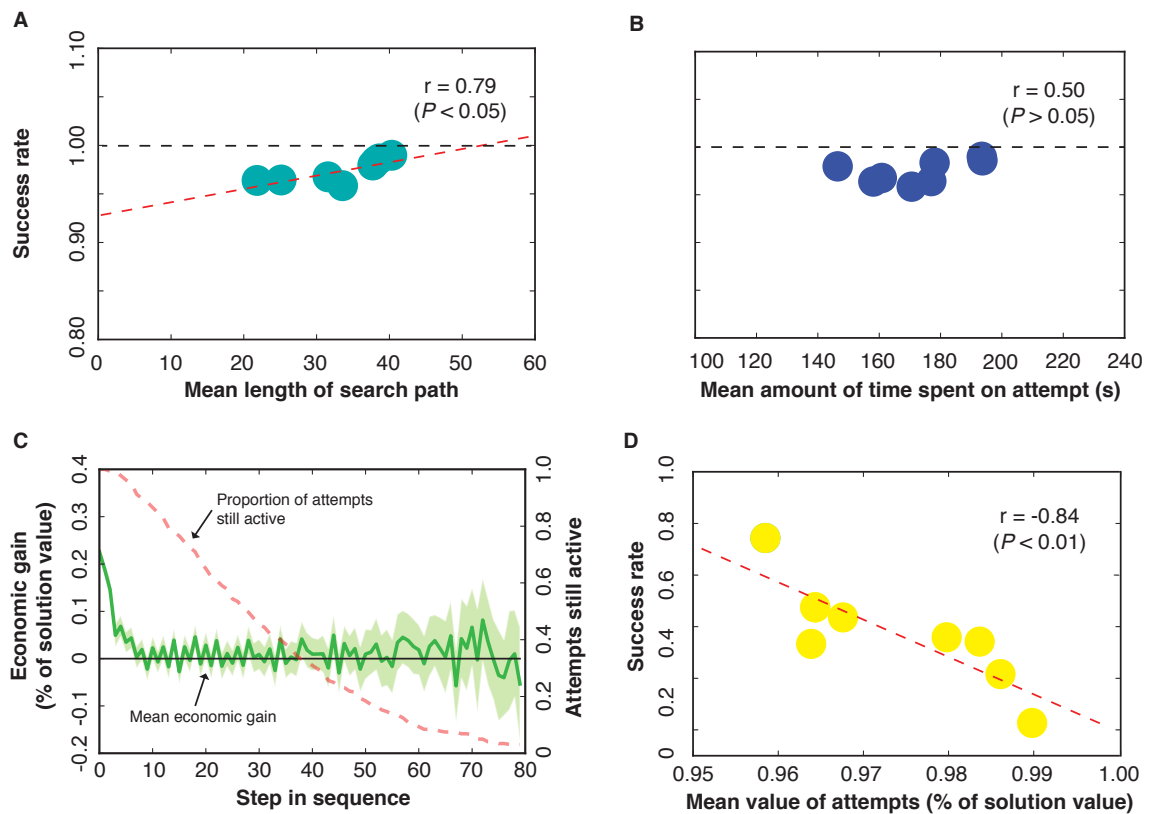
126 For a Turing machine, endowed with unlimited memory capacity, effort is measured as



127 number of computations. Therefore, we first measured effort as the path length in an attempt.  
128 This number can be considered as a proxy of the number of computations performed by the  
129 participant during an attempt, that is, a measure of computational time (analogous to *CPU time*  
130 in computing). We found that economic performance was not related to path length in attempts  
131 ( $P > 0.05$ , main effect of path length, linear mixed model (LMM) with random effects on  
132 intercept for individual instances and participants; SI Results 2.4, Table S3 Model 1; Fig. 3A).  
133 Path length was not related to computational performance, either ( $P > 0.05$ , main effect of path  
134 length, generalised linear mixed model (GLMM) with random effects on intercept for individual  
135 instances and participants; SI Results 2.4, Table S3 Model 3).

136 In principle, there was no opportunity cost to time in our experimental setting, because  
137 participants were barred from other activities while engaged in our task. Thus, clock time in  
138 itself cannot measure effort. However, clock time increases in the number of computations per-  
139 formed, as well as other dimensions of effort. The latter include memory retrieval, differences  
140 in effort required across types of computations, and other operations. These dimensions of ef-  
141 fort are not recognised as effort in Turing machines, but may constitute effort for humans. As  
142 such, clock time may provide a catch-all measure of effort. Indeed, we found that more time  
143 spent on solving an instance was related to higher economic performance ( $P < 0.05$ , main  
144 effect of clock time, LMM with with random effects on intercept for individual instances and  
145 participants; SI Results 2.4, Table S3 Model 2; Fig. 3B). Time spent on an attempt was not re-  
146 lated to computational performance ( $P > 0.05$ , main effect of clock time, GLMM with random  
147 effects on intercept for individual instances and participants; SI Results 2.4, Table S3 Model 4)

148 A traditional economic agent (“Homo economicus”) would search for a solution until  
149 marginal gain equaled marginal cost of effort. Using path length, a proxy for the number of



**Figure 3.** Performance and complexity. (A) Scatter plot of success rates and mean length of the search path in each of the instances. Black dashed line: 100%, red dashed line: regression line fitted to the data points,  $r$ : Pearson correlation between mean success rates and mean length of search path. (B) Scatter plot of success rates and mean amount of clock time spent on an attempt in each of the instances. (C) Mean value gain (in % of solution value) per step in the search path. (D) Scatter plot of success rates and mean value of attempts (normalised by solution value) in each of the instances.

150 computations, as a measure of effort, we plotted gain in economic performance against path  
151 length (Fig. 3C). Mean gain quickly dropped to zero but participants continued to search be-  
152 yond this point. Indeed, most of the search time was spent at zero net marginal gain. The  
153 same picture emerged when we plotted economic gain against time, which, as discussed above,  
154 provides a more general measure of effort (Fig. 3D, Fig. S2). Thus, participants violated one  
155 of the basic principles of economic theory. The same conclusion could be drawn if instead  
156 we assumed that participants were boundedly rational economic agents who worked until they  
157 reached an aspiration point (18). In that case, economic gains should have remained positive  
158 until they stopped spending effort, while they were not. One might argue that our results could  
159 be explained by unrealistically high aspiration levels, and that participants therefore continued  
160 to spend effort even if economic performance did not increase. However, in this case, we would  
161 expect participants to spend effort until the end of allotted time, which they almost never did.

162       If neither rational choice theory nor bounded rationality could explain economic perfor-  
163 mance, what did? We found that economic performance could be explained by instance dif-  
164 ficulty: Participants generated higher economic value on average in instances that were more  
165 difficult (Pearson correlation  $r = 0.838$ ;  $P < 0.01$ ; Fig. 3D). Here, difficulty was measured in  
166 terms of computational performance, assuming, quite reasonably, that if the score is lower, the  
167 instance was more difficult. The finding that economic performance increased with instance dif-  
168 ficulty is paradoxical from the point of view of economic theory. It also raises a question: How  
169 could participants have sensed which instances were more difficult? The KP is hard precisely  
170 because the only way to be sure that one has found the solution is to list *all* possible knapsacks  
171 with their corresponding values (2).

172 **Performance and computational complexity.** We investigated what constituted difficulty for  
173 participants, in particular, whether difficulty was related to computational complexity. The con-  
174 cept of computational complexity is typically defined based on the number of computations and  
175 the memory required by algorithms that find the solution of the problem (5). We considered as  
176 measures of complexity of an instance the variables that determine the number of computations  
177 and memory requirements for the major algorithms developed for the 0-1 knapsack problem (SI  
178 Materials and Methods 1.3, 1.5).

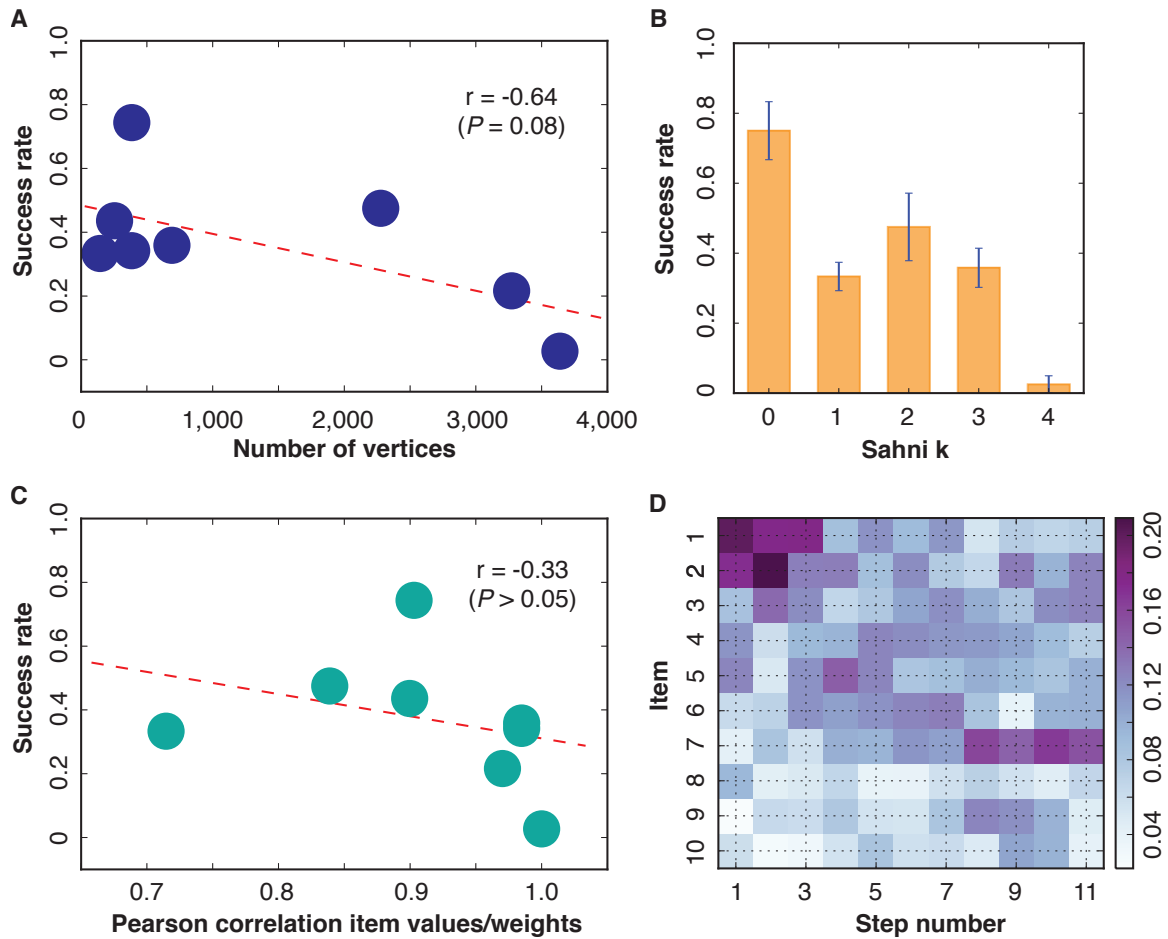
179 One class of algorithms are *uninformed* search algorithms, which find the solution by  
180 exhaustive search through the search space (19). The latter can be formalised as a graph and thus  
181 its size can be expressed as the number of vertices in this graph (SI Materials and Methods 1.2).  
182 We found that difficulty was negatively related to the total number of vertices in an instance's  
183 graph ( $P < 0.001$ , main effect of number of vertices, GLMM with participant random effects  
184 on intercept; SI Results 2.5, Table S4 Model 1), as well as to the number of terminal vertices  
185 ( $P < 0.01$ , main effect of number of terminal vertices, GLMM with participant random effects  
186 on intercept; SI Results 2.5, Table S4 Model 2).

187 Another class of algorithms are *informed* search algorithms, which use a set of rules or  
188 heuristics to guide search (19). The KP can be solved by dynamic programming (2). The  
189 dynamic programming algorithm represents the KP in a way that a Turing machine can solve  
190 the instance in polynomial time as a function of the size of the representation (SI Materials and  
191 Methods 1.3, 1.5). The representation is a two-dimensional matrix of size equal to (number of  
192 items)  $\times$  (base-2 logarithm of knapsack capacity), which is referred to as *input size*. For our first  
193 instance, for example, a Turing machine would need a memory of size  $10 \times \log_2 1900 \approx 109$   
194 bits, which is much larger than the capacity of human working memory (17). Thus, one would

195 not expect that humans were able to implement the dynamic programming algorithm in our  
196 experiment. Indeed, if one assumes that the number of computational mistakes increases in the  
197 number of computations required, and humans apply dynamic programming, one would predict  
198 that computational performance would decrease in input size. We found that computational  
199 performance was not related to input size ( $P > 0.05$ , main effect of input size, GLMM with  
200 participant random effects on intercept; SI Results 2.5, Table S4 Model 3).

201 An important algorithm for the 0-1 knapsack problem is the Sahni- $k$  algorithm (20) (SI  
202 Materials and Methods 1.3, 1.5). It solves an instance by a combination of brute-force search  
203 through a subset of items of cardinality  $k$ , and the greedy algorithm (21), filling up the knapsack  
204 by picking items in reverse order of their value-to-weight ratio. If  $k$  equals zero, the Sahni- $k$   
205 algorithm is similar to the greedy algorithm. If  $k$  is equal to the number of items in the solution,  
206 the algorithm is similar to a brute-force search through the entire search space.  $k$  is proportional  
207 to the number of computational steps and the memory required to compute the solution of an  
208 instance. Complexity of an instance can be defined as the minimum level of  $k$  necessary to  
209 solve the instance with the Sahni- $k$  algorithm. We found that computational performance was  
210 negatively correlated with Sahni- $k$  ( $P < 0.001$ , main effect of Sahni- $k$ , GLMM with participant  
211 random effects on intercept; SI Results 2.5, Table S4 Model 4; Fig. 4A).

212 The finding that Sahni- $k$  was related to difficulty suggests that the search algorithm par-  
213 ticipants were using was of lower complexity. Given that the score was highest for the instance  
214 with a Sahni- $k$  of zero, we hypothesised that participants were using an algorithm similar to the  
215 greedy algorithm, at least in the early stages of their attempts. We found that in the first few steps  
216 of their attempts, participants were more likely to add items with the highest value-to-weight  
217 ratio (Fig. 4B, Fig. S3). After a number of steps, the picture is less clear. Given that implemen-



**Figure 4.** Properties of search. (A) Scatter plot of number of vertices in instance graphs and success rates in instances. Red dashed line: regression line fitted to data,  $r$ : Pearson correlation between number of vertices and success rates. (B) Success rates in instances by Sahni- $k$  (see Materials and Methods). Blue lines: standard errors of means. (C) Scatter plot of Pearson correlation of item values and weights in an instance and success rate. (D) Time course of choice frequencies of items. The items available in an instance were sorted in reverse order of their density (value-to-weight ratio). The heat map shows average choice frequencies for the items for the first 11 steps in the search path across all attempts.


218 tation of the greedy algorithm depends on the correct computation of value-to-weight ratios,  
219 we would expect difficulty to increase in the correlation between item values and weights. This  
220 correlation is the so-called “conditioning number” of instances. When this number is high, even  
221 real-world computers would have difficulty finding the solution with any procedure involving  
222 the greedy algorithm (22). Nevertheless, in our case, while computational performance was  
223 negatively related to the conditioning number ( $P < 0.05$ , main effect of conditioning number,  
224 GLMM with participant random effects on intercept; SI Results 2.5, Table S4 Model 5), the  
225 conditioning number did not explain differences in computational performance across instances  
226 that Sahni- $k$  could not capture ( $P > 0.05$ , interaction effect Sahni- $k \times$  conditioning number,  
227 GLMM with participant random effects on intercept; SI Results 2.5, Table S4 Model 6). Pre-  
228 sumably this is because the greedy algorithm could find the solution only for one instance, and  
229 hence, in the other instances, inability to compute value-to-weight ratios would have been a  
230 virtue: one would have been forced to drop the greedy algorithm and consider alternative item  
231 sets.

232 We also found that participants searched longer and more successfully than the greedy al-  
233 gorithm (SI Results 2.7). Therefore, we conjectured that participants used a branch-and-bound  
234 algorithm, a type of algorithm that starts with the greedy algorithm but then searches for im-  
235 provements until a termination criterion is reached (23) (SI Materials and Methods 1.3). Since  
236 Sahni- $k$  is a measure of deviation of a solution from the greedy algorithm, we hypothesized that  
237 participants who tried to replace multiple items at once would realise a higher score, at least  
238 for instances where this was needed, that is, for high Sahni- $k$  instances. Therefore, we tested  
239 whether computational performance could be explained by the interaction between Sahni- $k$   
240 and the number of items participants removed at once on average after reaching the first full

241 knapsack. We measured the latter as the length of the shortest path between two full knapsack  
242 attempts, that is, between two terminal vertices (SI Materials and Methods 1.2). The interac-  
243 tion term was indeed significant ( $P < 0.01$ , interaction Sahni- $k \times$  mean distance, GLMM with  
244 random effect for instances on intercept; SI Results 2.8, Table S6 Model 1). As such, hetero-  
245 geneity in search strategies across participants could be partly explained by how many items  
246 participants reconsidered and changed during their searches.

247 In branch-and-bound search, each item is considered in turn, and an upper bound (some-  
248 times also a lower bound) to the value of the knapsack is computed, either with or without the  
249 item at hand. If this upper bound is higher with the item, then it is retained. Once retained,  
250 it will never be deleted. The latter induces inflexibility that makes the procedure fail (to find  
251 the optimum). This approach is useful when the decision-maker has episodic-memory con-  
252 straints, in the sense that he cannot remember what made him add the item earlier. An item  
253 may look optimal at the time of inclusion, and since the decision-maker cannot recall the entire  
254 search path, it is better to retain the item. We examined to what extent there was a tendency  
255 not to eliminate incorrect items that were added early on, and whether this determined compu-  
256 tational performance. To this end, we considered the distribution of the age of incorrect items  
257 that were eventually deleted (Fig. 5). We measured age as a fraction of number of steps taken  
258 since the beginning of an attempt (age equals 1 if the item was the first added to the knapsack).  
259 The vast majority of deleted incorrect items were added very recently ( $M = 0.2920$ ,  $SEM =$   
260  $0.0001$ ); only rarely did participants eliminate incorrect items that were added to the knapsack  
261 early on. A similar pattern emerged for correct items that were deleted ( $M = 0.2352$ ,  $SEM =$   
262  $0.0001$ ; Fig. 5). Mean age of correct items was significantly higher than age of incorrect items  
263 ( $P < 0.001$ , two-sample  $t$ -test,  $t = 6.98$ ) and their distributions were significantly different





**Figure 5.** Distribution of item ages. Histogram of age of items (SI Results 2.8). Age of an item was calculated as the fraction of number of steps taken since the beginning of an attempt (age equals 1 if the item was the first added to the knapsack). Green bars: correct items, red bars: incorrect items:  $M_{corr}$ : mean age and standard error of mean of correct items,  $M_{incorr}$ : mean age and standard error of mean of incorrect items.

264 ( $P < 0.001$ , Kolmogorov-Smirnov test for independence of samples,  $D = 0.10$ ). This means  
 265 that participants were reluctant to re-consider items that were added early on, introducing path  
 266 dependence in search. We also found that the initial full knapsack heavily determined eventual  
 267 computational success: whether the solution was found in an attempt depended significantly on  
 268 distance (in number of steps) to the solution of the first terminal vertex ( $P < 0.001$ , main effect  
 269 of distance of first vertex, LMM with random effects for participants and instances on intercept  
 270 and; SI Results 2.8, Table S7 Model 1).

271 We also examined variation in searches between participants. We found that there was  
 272 little overlap in the knapsacks that participants attempted. On average, individuals only visited

273 3.6% of all admissible knapsacks (Fig. S4A), yet combining all attempts, the twenty participants  
274 visited 42.1% of all admissible knapsacks (Fig. S4C, SI Results 2.6). This means that the group  
275 of 20 participants together explored more than fifteen times more of the search space than each  
276 individual participant explored on their own. This suggests that there was substantial hetero-  
277 geneity in search strategies (Video S1–16). However, in all instances, at least one participant  
278 found the solution quickly (Fig. S5).

279 The fact that the KP is NP-complete implies that determining the decision problem requires  
280 listing all possible solutions (SI Materials and Methods 1.1). In the theory of computation, the  
281 problem of finding the solution, referred to as the computational problem, is distinguished from  
282 the problem of deciding whether a given candidate solution is the solution of the problem,  
283 referred to as the decision problem. So far, we have considered the computational problem. To  
284 investigate whether participants solved the decision problem, that is, whether they knew that  
285 they found the solution, we examined subsequent attempts of the same instance. If a participant  
286 found the solution in the first attempt, and they solved the decision problem, then we would  
287 expect them to also solve the instance in the second attempt. We found that of those participants  
288 who found the solution in the first attempt, 31.2% did not solve the instance in the second  
289 attempt (SI Results 2.9), which suggests that participants often did not realise that they found  
290 the solution.

## 291 **Discussion**

292 Our results indicate that computational theory based on the concept of a Turing machine, as  
293 well as algorithms developed for real-world computers can help understand how humans solve

294 the KP (24–26). Like for Turing machines, we discovered that effort can be measured in terms  
295 of number of computations, but that this gives an incomplete picture. The finding that some par-  
296 ticipants improved computational performance because they were able to replace several items  
297 at once, while others could not, suggest that working memory is a binding constraint. Indeed,  
298 the replacement operation requires one to hold in memory several items at once (items deleted  
299 and items to be added). We also found that participants were reluctant to re-consider items that  
300 were added early on, thus causing eventual computational performance to depend heavily on the  
301 initial full knapsack. Unwillingness to re-consider ‘aged’ items suggest that episodic memory  
302 may be failing: the circumstances that led to initial inclusion are forgotten by the time the item  
303 is to be re-considered, and hence, the participant, lacking trust in episodic memory, decides  
304 not to entertain exclusion. Both findings point to the need for cognitive flexibility. One can  
305 easily imagine intervention policies that would enforce cognitive flexibility. It is an interesting  
306 question whether higher cognitive flexibility would indeed improve performance in the KP.

307       We discovered that various measures of complexity explained computational performance,  
308 in particular, the Sahni- $k$  measure. This same measure also explained computational perfor-  
309 mance in an earlier study in a group (markets) setting (13). The latter study also demonstrated  
310 that sharing of information through a market mechanism benefits individual decision-making.  
311 Our present findings explain why: there was little overlap between participants in the knapsacks  
312 that they visited, so that information sharing had the potential to vastly improve performance.  
313 To the extent that intellectual discovery and innovation can be thought of as solving instances  
314 of the KP (11–15), our results therefore question the effectiveness of an incentive system that  
315 discourages people from sharing information, such as the current patent system. Our earlier  
316 findings (13) confirm the relevance of this question.

317 We also found that computational performance decreased with the size of the search space  
318 measured by the size of the graph of an instance. This result may potentially explain earlier  
319 findings that the number of decision options negatively affects choice, a phenomenon referred  
320 to as “choice overload” (27). Our findings offer an algorithmic explanation of this phenomenon:  
321 the larger the number of choice options, the larger the search space through which the decision-  
322 maker has to search, and the less likely the decision-maker is able to discover the solution.

323 One of our key findings is that participants tended to spend more effort, and hence, en-  
324 hanced economic performance, on more difficult problems. This rather paradoxical result has  
325 profound implications for economic theory. According to the basic principles of economic the-  
326 ory, participants would be expected to spend effort until performance gain drops to marginal cost  
327 of effort, or until an aspiration point is reached. Our results may explain important situations  
328 where humans do not trade off marginal utility against marginal effort, such as in labour-leisure  
329 choices. For instance, taxi drivers do not increase time spent on the job when marginal gains are  
330 high; instead, they increase effort per unit of economic gain on “difficult” days (days when pas-  
331 sengers are harder to locate) (28). An important question is how participants in our experiment  
332 were able to recognise that an instance was more difficult, because the measures that explained  
333 difficulty cannot be constructed until one knows the entire search space.

334 Our results also cast doubt on another basic principle of economic theory, the assumption  
335 that decision-makers always optimise (29). Economic theory posits that in a given decision  
336 situation, decision-makers will always choose the option they prefer over all other available  
337 options, irrespective of the complexity of the decision problem. Our results suggest, however,  
338 that humans are unable to solve many, if not most, optimisation problems. Indeed, our findings  
339 render many current economic models of behaviour, including many models based on revealed

340 preferences theory, and the inferences drawn from them, implausible. Future developments of  
341 theories of choice need to take into account the computational demands that implementation of  
342 decisions requires (30–32).

343 It has been argued that decision-makers use heuristics to solve complex problems (33).  
344 This raises hitherto unanswered questions: (i) What is complexity as experienced by humans?,  
345 and (ii) Are the heuristics, and hence, performance, adapted to this complexity? Our results  
346 indicate that performance in the KP was driven by complexity of the instances, and that com-  
347 plexity resembled the notion of computational complexity, developed based on the concept of  
348 a Turing machine, when adapted suitably to accommodate constraints imposed by human bio-  
349 logical reality: limited computational capacity, working memory, and episodic memory.

350 Given the central importance of the KP in human life, a better understanding of how  
351 humans approach the KP is likely to lead to a more comprehensive understanding of human  
352 decision-making. It is bound to improve welfare not only through better incentivisation of  
353 intellectual discovery and innovation. It also promises to contribute to health through better di-  
354 agnosis and treatment of mental disorders that implicate choice that is detrimental to successful  
355 KP solving, because of lack of cognitive flexibility (obsessive-compulsive disorder (34)), sub-  
356 optimal time allocation (attention deficit/hyperactivity syndrome (35)), or stringent adherence  
357 to “rational” economic principles (autism spectrum disorder (36)).

## 358 **Materials and Methods**

359 **Participants and experimental task.** Twenty human volunteers (age range = 18–30, mean  
360 age = 21.9, 10 female, 10 male), recruited from the general population, took part in the study.

361 Inclusion criteria were based on age (minimum = 18 years, maximum = 35 years), right-  
362 handedness and normal or corrected-to-normal vision. The experimental protocol was approved  
363 by The University of Melbourne Human Research Ethics Committee (Ethics ID 1443290), and  
364 written informed consent was obtained from all participants prior to commencement of the ex-  
365 perimental sessions.

366 Participants were asked to solve eight instances of the 0-1 knapsack problem (2). For  
367 each instance, participants had to select from a set of items of given values and weights, the  
368 subset of items with the highest total value, subject to a total weight constraint (Table S1). The  
369 instances used in this study were used in a prior study (13) and differed significantly in their  
370 computational complexity (Table S2).

371 The instances were displayed on a computer display (1000 x 720 pixels; Fig. 1B). Each  
372 item was represented by a square. Value and weight of an item were displayed at the centre of  
373 the square. The size of an item was proportional to its weight and the colour (share of blue)  
374 was proportional to its value. At the top of the screen, total value, total weight and weight  
375 constraint of the knapsack were displayed. When the mouse was moved over an item, a black  
376 frame around the square appeared and the counters at the top of the screen added this items'  
377 value and weight to the totals. When the mouse was moved over an item that could not be added  
378 to the knapsack at that time, because its addition would have violated the weight constraint, the  
379 counters turned red. An item was selected into the knapsack by clicking on it. Once an item  
380 was selected into the knapsack, it turned green. The item stayed green until it was removed  
381 from the knapsack (by clicking on it again). A solution was submitted by pressing the space  
382 bar. An attempt was automatically terminated after 240 s and time remaining was displayed by  
383 a progress bar in the top-right corner of the screen.

384 Each participant had two attempts per instance. The order of instances was randomised  
385 across an experimental session. We recorded the time course of selection of items to and re-  
386 movals from the knapsack. To make the task incentive compatible, participants received a pay-  
387 ment proportional to the values of their attempts (between \$0 and \$4 per attempt). In addition,  
388 participants received a show-up fee of \$5.

389 **Data analysis.** For each attempt, we recorded the sequence of additions of items to and re-  
390 movals of items from the knapsack. Each element in this sequence represents a state of the  
391 knapsack, and each state of the knapsack corresponds to a vertex in the graph  $G$  of the instance  
392 (the first element of the sequence always corresponds to the initial vertex of  $G$ , and the last  
393 element always corresponds to the participant’s proposed solution of the instance). A sequence  
394 of additions and removals can be represented as a path in the graph (see SI Materials and Meth-  
395 ods 1.2).

396 For each attempt, we recorded the time when the attempt was submitted as well as the  
397 sequence of additions and removals of items. For each step in this sequence, we computed the  
398 total value of items selected as well as the distance  $\ell_G(i, s)$  to the solution vertex  $s$  from the  
399 vertex  $i$  in the graph representing this subset of items (see SI Materials and Methods 1.2). The  
400 subset of items selected at the time of submission was the participant’s proposed solution of the  
401 instance. The attempt was marked correct if the subset of items in the participant’s proposed  
402 solution was the solution of the instance (that is,  $\ell_G(i, s)$  was equal to zero), and incorrect  
403 otherwise.

404 To evaluate an attempt in value space, we computed the value of the proposed solution  
405 normalised by the value of the solution, which corresponds to the reward schedule. We also

406 computed the difference between the proposed solution and the mean of the values of all termi-  
407 nal vertices in the graph representing the problem. The latter is the mean of the values of all  
408 maximally admissible knapsacks, which is equal to the expected value of randomly selecting  
409 items into the knapsack until the knapsack is full.

410 All analyses were performed in Python (version 2.7.6) and R (version 3.2.0).

411

412 **Author contributions:** P.B. and C.M. designed research, performed research, analyzed data,  
413 and wrote the manuscript.

414

415 **Acknowledgements:** The authors are grateful to Haley Warren for support with data acquisition  
416 and to Debrah Meloso for constructive feedback on the manuscript. C.M. was supported by a  
417 Strategic Research Initiatives Fund grant from the Faculty of Business and Economics, The  
418 University of Melbourne. The authors declare that they have no competing financial interests.

## 419 **References**

420 [1] Mathews GB (1897) On the partition of numbers. *Proceedings of the London Mathemati-*  
421 *cal Society* 28:486–490.

422 [2] Kellerer H, Pferschy U, Pisinger D (2004) *Knapsack Problems*. (Springer Science &  
423 Business Media).

424 [3] Cook SA (1983) An overview of computational complexity. *Communications of the ACM*  
425 26(6):400–408.



- 426 [4] Karp RM (1972) Reducibility among Combinatorial Problems in *Complexity of Computer*  
427 *Computations*. (Springer US), pp. 85–103.
- 428 [5] Moore C, Mertens S (2011) *The Nature of Computation*. (Oxford University Press, Ox-  
429 ford).
- 430 [6] Sims CA (2006) Rational inattention: Beyond the linear-quadratic case. *The American*  
431 *Economic Review* 96(2):158–163.
- 432 [7] Torralva T, Gleichgerrcht E, Lischinsky A, Roca M, Manes F (2013) “Ecological” and  
433 highly demanding executive tasks detect real-life deficits in high-functioning adult ADHD  
434 patients. 17(1):11–19.
- 435 [8] de Carvalho V (1998) Exact solution of cutting stock problems using column generation  
436 and branch-and-bound. *International Transactions in Operational Research* 5:35–44.
- 437 [9] Pisinger D (2002) Heuristics for the container loading problem. *European Journal of*  
438 *Operational Research* 141(2):382–392.
- 439 [10] Mansini R, Speranza MG (1999) Heuristic algorithms for the portfolio selection problem  
440 with minimum transaction lots. *European Journal of Operational Research* 114(2):219–  
441 233.
- 442 [11] Arthur WB, Polak W (2006) The evolution of technology within a simple computer model.  
443 *Complexity* 11:23–31.
- 444 [12] Arthur WB (2009) *The Nature of Technology: What It Is and How It Evolves*. (Allen Lane,  
445 London).

- 446 [13] Meloso D, Copic J, Bossaerts P (2009) Promoting intellectual discovery: patents versus  
447 markets. *Science* 323(5919):1335–1339.
- 448 [14] Wagner A (2014) *Arrival of the Fittest: Solving Evolution’s Greatest Puzzle*. (Oneworld,  
449 London).
- 450 [15] Youn H, Strumsky D, Bettencourt LMA, Lobo J (2015) Invention as a combinatorial pro-  
451 cess: evidence from us patents. *Journal of the Royal Society Interface* 12:20150272.
- 452 [16] Fortnow L (2009) The status of the P versus NP problem. *Communications of the ACM*  
453 52(9):78–86.
- 454 [17] Cowan N (2010) The magical mystery four: How is working memory capacity limited,  
455 and why? *Current Directions in Psychological Science* 19(1):51–57.
- 456 [18] Simon HA (1959) Theories of decision-making in economics and behavioral science. *The*  
457 *American Economic Review* pp. 253–283.
- 458 [19] Russell S, Norvig P (2014) *Artificial Intelligence*. (Pearson, Harlow).
- 459 [20] Sahni S (1975) Approximate algorithms for the 0-1 knapsack problem. *Journal of the*  
460 *ACM* 22:115–124.
- 461 [21] Cormen TH, Leiserson C, Rivest RL, Stein C (2001) *Introduction to Algorithms*. (MIT  
462 Press, Cambridge, MA).
- 463 [22] Pisinger D (2004) Where are the hard knapsack problems? *Computers & Operations*  
464 *Research* 32:2271–2284.

- 465 [23] Land AH, Doig A (1960) An automatic method for solving discrete optimization prob-  
466 lems. *Econometrica* 28(3):497–520.
- 467 [24] Newell A, Simon H (1972) *Human Problem Solving*. (Prentice-Hall, Englewood Cliffs).
- 468 [25] Marr D, Poggio T (1976) From understanding computation to understanding neural cir-  
469 cuitry.
- 470 [26] Marr D (1982) *Vision*. (W.H. Freeman, San Francisco).
- 471 [27] Toffler A (1970) *Future Shock*. (Random House, New York, NY).
- 472 [28] Camerer C, Babcock L, Loewenstein G, Thaler R (1997) Labor supply of new york city  
473 cabdrivers: One day at a time. *The Quarterly Journal of Economics* pp. 407–441.
- 474 [29] McFadden D (1999) Rationality for economists? *Journal of Risk and Uncertainty*  
475 19(1):73–105.
- 476 [30] Lewis RL, Howes A, Singh S (2014) Computational rationality: Linking mechanisms and  
477 behavior through bounded utility maximisation. *Trends in Cognitive Science* 6:279–311.
- 478 [31] Gershman SJ, Horvitz EJ, Tenenbaum JB (2015) Computational rationality: A converging  
479 paradigm for intelligence in brains, minds, and machines. *Science* 349(6245):273–278.
- 480 [32] Parkes DC, Wellman MP (2015) Economic reasoning and artificial intelligence. *Science*  
481 349(6245):267–272.
- 482 [33] Tversky A, Kahneman D (1974) Judgment under Uncertainty: Heuristics and Biases. *Sci-*  
483 *ence* 185(4157):1124–1131.

- 484 [34] Clayton IC, Richards JC, Edwards CJ (1999) Selective attention in obsessive–compulsive  
485 disorder. *Journal of Abnormal Psychology* 108(1):171.
- 486 [35] Adler LA, Chua HC (2002) Management of ADHD in adults. *Journal of Clinical Psychi-*  
487 *atry* 63:29–35.
- 488 [36] Sally D, Hill E (2006) The development of interpersonal strategy: Autism, theory-of-mind,  
489 cooperation and fairness. *Journal of Economic Psychology* 27(1):73–97.

# 9 SUPPORTING INFORMATION

## 10 1 Materials and Methods

### 11 1.1 The 0-1 knapsack problem

12 The 0-1 knapsack problem is the problem of finding in a set of items of given values and  
13 weights the subset of items with the highest total value, subject to a total weight constraint  
14 (1, 2). Mathematically, the problem can be written as

$$\max \sum_{i=1}^I v_i w_i \text{ subject to } \sum_{i=1}^I w_i x_i \leq C \text{ and } x_i \in \{0, 1\}, \quad (1)$$

15 where  $I$  is the total number of items,  $v_i$  and  $w_i$ ,  $i = 1, \dots, I$ , denote value and weight, respec-  
16 tively, of an item and  $C$  is the capacity (maximum weight) of the knapsack.

17 The 0-1 knapsack problem is a combinatorial optimisation problem. It is a member of  
18 the complexity class *non-deterministic polynomial-time (NP) complete* (2–4). In computer sci-  
19 ence, complexity classes indicate the ‘hardness’ of computational problems. The members  
20 of these classes differ in the rate at which the number of computational steps (time complex-  
21 ity) or memory requirements (space complexity) grow as the size of a problem’s instance in-  
22 creases. An important distinction is made between problems for which computational time  
23 increases as a polynomial of the problem’s size (class polynomial-time, or P), and those prob-  
24 lems for which computational time increases faster than polynomially (class non-deterministic  
25 polynomial-time, or NP). If an algorithm exists that solves a problem in polynomial time, it is  
26 called “efficient” (5). Thus, members of class P are those problems for which there exist effi-  
27 cient solution algorithms, whereas for members of class NP, no efficient algorithms are known.

28 For those problems, verifying that one has found the solution of an instance requires listing the  
29 values of all possible solutions (brute-force search). This feature obtains despite the fact that  
30 one can compute relatively fast (time only increases as a polynomial of the number of items)  
31 whether a given solution reaches a certain value level. Membership of a complexity class is  
32 determined based on the hardest instances of a problem and some instances of a given problem  
33 may require less time and memory than others. At present, it is still unknown which properties  
34 of computational problems determine their complexity, referred to as the “P versus NP prob-  
35 lem”, which is one of the fundamental mathematical problems of our time (5). NP-complete  
36 problems have the additional property that if there is an efficient algorithm for one member of  
37 their class, then there exists an efficient algorithm for all members of this class as well as any  
38 problem in class NP (5).

## 39 **1.2 Representing instances of the knapsack problem as graphs**

40 An instance of the 0-1 knapsack problem can be represented as an undirected graph  $G = (V, E)$   
41 comprising vertices,  $V$ , and edges,  $E$  (6). We call a subset of items (knapsack) *admissible* if  
42 the combined weight of the items is less than or equal to the capacity of the knapsack,  $C$ . Each  
43 admissible subset  $s$  of items is represented by a vertex  $i \in V$ . We define the *order* of a graph  
44  $|G|$  as the number of vertices of the graph. Note that  $|G|$  will usually be lower than the number  
45 of all possible subsets of items, which is equal to  $2^I + 1$  (including the empty set), because  
46 some possible subsets are not admissible due to the weight constraint. We define value  $v_i$  and  
47 weight  $w_i$  of a vertex  $i$  as the sum of the values and weights, respectively, of the subset of  
48 items represented by vertex  $i$ . Two vertices  $i, j \in V$  are connected by an edge  $(i, j)$  if vertex  
49  $j$  can be reached from vertex  $i$  by adding one item to or removing one item from the knapsack

50 represented by vertex  $i$ , that is, if the difference between the sets of items represented by the  
 51 two vertices contains exactly one item. Because the graph  $G$  is undirected,  $(i, j) = (j, i)$  for all  
 52  $i, j \in V$ . We call a vertex  $i$  *incident* with edge  $e$  if  $i \in e$ , and we call the two vertices  $i, j \in G$   
 53 connected by edge  $(i, j)$  *adjacent* to each other. We define the *degree*  $d_G(i)$  of vertex  $i$  as the  
 54 number  $|E(i)|$  of edges at  $i$ . We assign each edge  $(i, j) \in E$  a weight  $w_{ij}$  equal to 1. A *path* is  
 55 a graph  $P = (V', E')$  of the form  $V = \{x_0, x_1, \dots, x_k\}$  and  $E = \{x_0x_1, x_1x_2, \dots, x_{k-1}x_k\}$ ,  
 56 where the  $x_i$  are all distinct. The vertices  $x_0$  and  $x_k$  are linked by  $P$ . We define the *length* of  
 57  $P$  as the sum of the weights of its edges. We define the *distance*  $\ell_G(i, j)$  in  $G$  of two vertices  
 58  $i, j$  as the length of a shortest  $i$ - $j$  path in  $G$ . The distance  $\ell_G$  is conceptually related to the *edit*  
 59 *distance* often used in computer science (7). We call the graph  $G$  representing a given instance  
 60 of the 0-1 knapsack problem the *graph induced by the instance*.

61 We also define the undirected graph  $\bar{G} = (\bar{V}, \bar{E})$  with vertices  $\bar{V}$  and edges  $\bar{E}$  defined as  
 62 in  $G$  except that edge weights are set equal to the value of the item whose addition to or removal  
 63 from the knapsack is represented by the edge. Paths in  $\bar{G}$  are defined similarly to paths in  $G$ .  
 64 The distance  $\ell_{\bar{G}}(i, j)$  of the two vertices  $i$  and  $j$  now represents the difference in values of the  
 65 two vertices. The solution of the knapsack problem represented by graph  $\bar{G}$  can be found by  
 66 computing the longest path in  $\bar{G}$  (8). We call the vertex  $i \in V$  representing the solution of the  
 67 instance the *solution node*.

68 Finally, we define the directed graph  $\vec{G} = (\vec{V}, \vec{E})$  with vertices  $\vec{V}$  as defined in  $G$ . Two  
 69 vertices  $i, j \in V$  are connected by an edge  $(i, j)$  if vertex  $j$  can be reached from vertex  $i$  by  
 70 adding one item to the knapsack represented by vertex  $i$ , and vice versa for removals of items.  
 71 Note that in graph  $\vec{G}$ ,  $(i, j) \neq (j, i)$  for all  $i, j \in \vec{V}$ . We assign each edge  $(i, j) \in \vec{E}$  a weight  
 72  $w_{ij} = v_j - v_i$ . Paths in  $\vec{G}$  are defined as in  $G$ . The distance  $\ell_{\vec{G}}(i, j)$  of the two vertices  $i$

73 and  $j$  represents the difference in values of the two vertices. We define the *out-degree*  $d_{\vec{G}}^{out}(i)$   
74 of vertex  $i$  as the number of edges leaving vertex  $i$ , and the *in-degree*  $d_{\vec{G}}^{in}(i)$  of vertex  $i$  as the  
75 number of edges terminating at vertex  $i$ . We call the vertex in  $\vec{G}$  representing the empty set  
76 (knapsack) the *initial vertex*. The initial vertex has an in-degree equal to zero. We call a vertex  
77 with out-degree equal to zero a *terminal vertex*. Each terminal vertex represents a *maximally*  
78 *admissible* knapsack, that is, a subset of items with the property that no additional remaining  
79 item could be added to the knapsack without violating the weight constraint. Note that the  
80 set of terminal vertices contains the vertex representing the solution of the knapsack problem.  
81 We consider all three graphs,  $G$ ,  $\bar{G}$  and  $\vec{G}$ , in the analysis of participants' attempts at solving  
82 knapsack problems.

83 Let us consider the graph  $G = (V, E)$  of some instance of the 0-1 knapsack problem.  
84 The economic value of each node is given by  $v_i$  for all  $i \in V$ . Let vertex  $i$  represent the  
85 initial node (empty knapsack) and vertex  $s$  represent the solution vertex. The distance between  
86 the two vertices  $\ell_G(i, s)$  is equal to number of items in the solution of the instance. We can  
87 compute  $\ell_G(i, s)$  for all other vertices  $i \in G, i \neq s$ . Intuitively, for any  $i \in G, i \neq s$ ,  $\ell_G(i, s)$   
88 equals the number of additions of items to and removals of items from the knapsack to get from  
89 the knapsack represented by vertex  $i$  to the solution of the instance, represented by vertex  $s$ .  
90 The mean correlation between vertex values  $v$  and their distances to the solution vertex  $\ell_G$  in  
91 the instances investigated in this study was  $-0.22$  (min =  $-0.41$ , max =  $0.04$ , SD =  $0.13$ ;  
92 Tab. S2). Note that in convex problems this correlation would be positive. The low correlation  
93 between values and distances of vertices is one aspect of the knapsack problem that makes it  
94 hard. It means that optimisation algorithms based on local increase in marginal value such as  
95 hill-climbing do not work for the knapsack problem in general.



96 To illustrate this property, we plot the graphs of the instances investigated in this study in  
97 value-item (distance) space (Fig. S1). The position of each node in the graph of an instance  
98 is determined by its value (normalised by the value of the solution; abscissa) and distance  
99 (ordinate). The initial node is indicated in yellow and the solution node is indicated in red (top-  
100 right corner). As the plots illustrate, in each of the instances, there are many vertices of equal  
101 value but different distances, and vice versa.

### 102 **1.3 Computational approaches to solving the 0-1 knapsack problem**

103 Various algorithms have been proposed for the 0-1 knapsack problem. An algorithm is a tool  
104 for solving a well-specified computational problem (9). It describes a specified computational  
105 procedure for achieving a desired relation between one set of values (input) and another set of  
106 values (output) that provides the solution to the computational problem. While every algorithm  
107 solves a particular computational problem, a given computational problem can often be solved  
108 by many different algorithms. This has led to the proposition that computational problems can  
109 be investigated separately from the algorithms that are used for solving these problems, that is,  
110 that the computational layer and the algorithmic layer are independent (10). More recently, it  
111 has been suggested that computational and algorithmic layers are often interdependent and that  
112 therefore the study of the algorithmic layer can often provide important insights into the nature  
113 of the computational problem. This is also relevant for economics because the discipline has  
114 traditionally focused on the characterisation of the computational problems agents solve and  
115 ignored the way in which agents solve these problems. We propose that studying the search  
116 algorithms that humans may have used may give us important clues about the optimisation  
117 problem they were trying to solve.

118 Two classes of algorithms for the 0-1 knapsack problem can be distinguished: *uninformed*  
119 and *informed* search algorithms. Uninformed algorithms, such as breadth-first or depth-first  
120 search (11), typically search the entire graph of an instance to find the solution. Alternatively,  
121 the solution of an instance represented by graph  $\bar{G}$  can also be found by computing the longest  
122 path in  $\bar{G}$  (8). Both, computational time and memory requirements of those algorithms increase  
123 non-polynomially in the size of the problem.

124 Informed search algorithms use some rule, sometimes referred to as *heuristic* (12), to guide  
125 the search. Instances of the 0-1 knapsack problem can be solved by dynamic programming (2).  
126 Here, the time to compute the solution of a given instance is proportional to the *input size* of the  
127 problem given by  $I \log_2 C$ , where  $I$  is the number of items in the instance and  $C$  is the capacity  
128 (weight limit) of the knapsack. Computing time and memory requirements of the dynamic  
129 programming algorithm still increase fast in the size of the problem and hence computation  
130 quickly becomes intractable.

131 Given the computational intractability of these approaches, various *approximation algo-*  
132 *rithms* have been developed for the knapsack problem. One important approximation algorithm  
133 is the *greedy algorithm* (9). It solves the knapsack problem by selecting items according to  
134 decreasing *density* of the items, where density is defined as the ratio of value to weight of an  
135 item. The greedy algorithm has much lower computational demands than dynamic program-  
136 ming. It only requires sorting of the items. However, it is not guaranteed to find the solution of  
137 an instance and its proposed solution may be arbitrarily far away from the solution. Note that  
138 the greedy algorithm always finds the optimum in a variation of the instance in which fractions  
139 of items are allowed, that is,  $w_i \in [0, 1], i = 1, \dots, I$  (LP-relaxation of the instance). The value  
140 of the solution in this modified problem (LP-bound) is often used as an approximation of the

141 value of the solution in the 0-1 knapsack problem.

142 Another important type of algorithm is the *branch-and-bound* algorithm (13). This algo-  
143 rithm starts with the greedy algorithm to construct an initial attempt and subsequently optimises  
144 the knapsack by selectively removing and adding items until a termination criterion has been  
145 reached. Like the greedy algorithm, the branch-and-bound algorithm is not guaranteed to find  
146 the solutions of all instances of the 0-1 knapsack problem.

147 We also consider the *Sahni-k* algorithm (14). The Sahni algorithm of order  $k$  considers a  
148 subset of  $k$  items and fills the knapsack using the greedy algorithm. It does so for all possible  
149 combinations of  $k$  items. The proposed solution is the knapsack with the highest total value.  
150 Note that the Sahni algorithm of order zero is equal to the greedy algorithm. If  $k$  is equal to the  
151 number of items in the solution of the instance, then the algorithm is equivalent to a brute force  
152 search and the solution proposed by the algorithm will be the solution of the instance. However,  
153 if  $k$  is less than the number of items in the solution, then the Sahni algorithm is not guaranteed  
154 to find the solution of the instance.

#### 155 **1.4 Example instance of the 0-1 knapsack problem**

156 We briefly discuss a small example instance of the 0-1 knapsack problem to illustrate the con-  
157 cepts discussed above. The properties of the instance are displayed in Fig. 1. There are five  
158 items available and the capacity of the knapsack is 7. Thus, the number of possible combina-  
159 tions of items is  $2^5 = 32$  but only 18 combinations are feasible, that is, they meet the weight  
160 constraint. The solution is the set (1, 3) with total value 21 and weight 7.

161 At the bottom of the Fig. 1, the graph of the instance is displayed. The positions of vertices  
162 in the x-y plane are determined by their value (x-axis) and distance to the solution node (y-

163 axis). The graph has 18 vertices, 31 edges connecting the vertices, and 7 terminal vertices (that  
164 is, maximally admissible sets). The empty set (initial vertex) is displayed in yellow and the  
165 solution vertex is displayed in red. The graph illustrates the low correlation between values  
166 and distance of the solution of the vertices. Some of the vertices have similar values but are at  
167 different distances to the solution, whereas other vertices have the same distance to the solution  
168 but different values.

169 As is typical of the 0-1 knapsack problem, some of the approximation algorithms would  
170 not find the solution of the instance. For example, the greedy algorithm would select items in  
171 decreasing order of their value-to-weight ratio and end up with set (1, 2, 4) with total value 20  
172 and weight 7.

## 173 **1.5 Measuring difficulty of an instance**

174 In computer science, problems are classified according to computational complexity. The most  
175 important classification is based on the time it takes a computer (Turing machine) to compute  
176 the solution of the problem, irrespective of the algorithm used. The problem of finding the  
177 solution of the 0-1 knapsack problem is a member of the complexity class NP-hard (2). For  
178 these classes, there is no known algorithm with the properties that the solution is correct and  
179 the time to compute the solution is a polynomial of the problem's size (number of items in the  
180 problem). The members of class NP-hard are considered the hardest computational problems  
181 currently known. The decision problem of determining whether a candidate solution is the  
182 solution of the problem is a member of the complexity class NP-complete (2). Membership of  
183 a complexity class is based on the hardest instances of a problem, that is, instances of a given  
184 problem may vary in difficulty.

185       Difficulty can also be expressed in terms of various measures of the topology of the graph  
186 induced by an instance. They include the number of vertices in the graph representing the  
187 problem (number of admissible sets) and the number of terminal vertices in the graph (number  
188 of maximally admissible sets). These measures can be regarded as properties of the search  
189 space through which algorithms have to search in order to find the solution, and computational  
190 time and memory requirements of most algorithms increase with the size of the search space.

191       The 0-1 knapsack problem is a special case of the class NP-hard because its instances can  
192 be solved by dynamic programming. The computational time of the dynamic programming  
193 algorithm for the 0-1 knapsack problem is proportional to  $I \log_2 C$ , where  $I$  is the number of  
194 items available in the instance and  $C$  is the capacity of the knapsack. Hence, the problem  
195 is said to be *pseudo-polynomial*. We will refer to  $I \log_2 C$  as the *input size* of the dynamic  
196 programming algorithm. It will be one of the measures of difficulty of instances.

197       Difficulty can also be expressed in terms of the complexity of a search algorithm. One  
198 measure we consider here is based on the Sahni- $k$  algorithm described above (14, 15). We  
199 define the  $k$  of a given instance as the smallest order of the Sahni- $k$  algorithm that finds the  
200 exact solution of the instance. For example, an instance with  $k = 0$  is an instance that can  
201 be solved with the Sahni algorithm of order 0, that is, the problem can be solved by applying  
202 the greedy algorithm. The higher the  $k$  of an instance, the larger the distance of the instance's  
203 exact solution from the solution computed by the greedy solution. The higher the value of  $k$ ,  
204 the higher are the number of computations and the memory requirement of the algorithm. The  
205 value of  $k$  of the instances considered in the present study ranges from 0 to 4 (Tab. S2).

206       Another measure of difficulty of instances of the 0-1 knapsack problem is the Pearson cor-  
207 relation coefficient between values and weights of the available items (16). Stronger correlation

208 is associated with greater difficulty. In many real-life situations, value and weight are strongly  
209 correlated. For example, in many investment problems, the return is proportional to the in-  
210 vestment outlay plus a fixed charge for each project. If value and weight of items are strongly  
211 correlated, the instance is hard to solve for two reasons. Firstly, there is a large gap between  
212 the continuous (LP relaxation) and integer solution of the problem, and thus the problem is  
213 ill-conditioned. Secondly, sorting the items in decreasing order of their value-to-weight ratio  
214 means sorting according to their weights. This means, though, that for any small interval of the  
215 ordered items, there is only limited variation in weights, making it more difficult to satisfy the  
216 capacity constraint with equality (16).

217       Importantly, the measures of computational complexity described above are all defined  
218 relative to a Turing machine, a mathematical model of a general computing machine. It is not  
219 clear to which extent they also apply to humans. However, we expect that many of the properties  
220 of instances that make them hard for a Turing machine also make them hard for humans. For  
221 example, the input size is proportional to the amount of memory required by the dynamic pro-  
222 gramming algorithm. Since human working memory is constrained, we would expect instances  
223 to become more difficult for humans as memory requirements increase, albeit at a different rate.  
224 Similarly, the Sahni- $k$  measure is a measure of the size of a combinatorial problem that needs  
225 to be solved in order to find the solution of an instance. For humans to solve combinatorial  
226 problems, they require working memory and they need to perform arithmetic. Thus, we expect  
227 that humans will perform worse on instances that require more memory, that is, instances with  
228 higher input size and higher Sahni- $k$  measures, *ceteris paribus*. On the other hand, in instances  
229 with low Sahni- $k$ , most items are selected based on the greedy algorithm, which requires sorting  
230 according to the value-to-weight ratio. Since humans are prone to mathematical mistakes, we

231 expect them to perform worse on instances with a high correlation between values and weights  
232 of items, *ceteris paribus*.

## 233 **2 Results**

### 234 **2.1 Duration of attempts**

235 In the following, we will only consider attempts that were submitted within the time limit of  
236 240 s. Of all 320 attempts in the experiment, 12 were not submitted within the limit, leaving  
237 308 attempts for analysis. The mean time spent on an attempt was 172.0 s ( $SD = 57.1$ ). Means  
238 of instances ( $min = 146.5$  s,  $M = 172.3$  s,  $max = 193.7$  s,  $SD = 15.7$  s) were not significantly  
239 different (one-way ANOVA,  $F_{(1,6)} = 5.2$ ,  $P = 0.06$ ). We also fitted survival functions sep-  
240 arately for each instance. We found that survival times differed significantly across instances  
241 (log-rank test,  $\chi^2_{(7)} = 14.9$ ,  $P < 0.05$ ). Participant means ( $min = 73.4$  s,  $M = 172.5$  s,  $max$   
242  $= 226.7$  s,  $SD = 39.8$  s) were not significantly different from each other (one-way ANOVA,  
243  $F_{(1,18)} = 0.13$ ,  $P > 0.05$ ) but survival times differed significantly across participants (log-rank  
244 test,  $\chi^2_{(19)} = 303$ ,  $P < 0.001$ ).

### 245 **2.2 Quality of attempts**

246 *Success rates:* The mean success rate, that is the proportion of attempts in which participants  
247 found the solution of an instance, was 37.4% ( $SD = 48.3\%$ ). In comparison, the expected  
248 success rate of an algorithm that fills knapsacks by picking items at random, which is equivalent  
249 to picking a maximally feasible knapsack at random, was 0.7%. The total number of successes

250 was significantly above chance (one-sided binomial test,  $P < 0.001$ ). The success rate varied  
251 substantially by both problem instance ( $min = 2.7\%$ ,  $M = 36.7\%$ ,  $max = 74.4\%$ ,  $SD = 19.3\%$ ;  
252 Fig. 2A) and participant ( $min = 6.2\%$ ,  $M = 37.4\%$ ,  $max = 56.2\%$ ,  $SD = 15.7\%$ ; Fig. 2B). One of  
253 the instances was only solved once and the participant who solved it had an overall success rate  
254 of 50.0% (there were 5 participants with higher average success rates). Note that performance  
255 varied more between problems (range = 71.7%) than between participants (range = 50.0%).

256 *Distance*: A refined measure of the quality of an attempt is the distance  $\ell_G$  of an attempt from  
257 the solution in the graph  $G$  induced by the instance (SI Materials and Methods 1.2). The mean  
258 distance was 2.639 ( $SD = 2.325$ ). It was significantly lower than the mean distance of attempts  
259 of an algorithm filling the knapsack by picking items at random, which was 5.068 (one-sample  
260 t-test,  $t_{(307)} = -18.374$ ,  $P < 0.001$ ). Distance, too, varied significantly by both instance ( $min$   
261  $= 0.784$ ,  $M = 2.622$ ,  $max = 4.865$ ,  $SD = 1.170$ ) and participant ( $min = 1.429$ ,  $M = 2.595$ ,  $max$   
262  $= 4.062$ ,  $SD = 0.765$ ).

263 *Economic value*: To assess economic performance, we computed the value of a participant's  
264 attempt and normalised it by the value of the solution. Mean economic performance was 97.4%  
265 ( $SD = 5.8\%$ ). It was significantly higher than the expected economic performance of an al-  
266 gorithm that fills knapsacks by randomly picking items until the knapsack is full, which was  
267 85.3% (one-sample t-test,  $t_{(307)} = 36.382$ ,  $P < 0.001$ ). Similar to the previous performance  
268 measures, economic performance varied more by instance ( $min = 95.8\%$ ,  $M = 97.4\%$ ,  $max =$   
269  $99.0\%$ ,  $SD = 1.1\%$ ) than by participant ( $min = 88.9\%$ ,  $M = 97.4\%$ ,  $max = 99.3\%$ ,  $SD = 2.4\%$ ).

270 A stricter criterion to assess economic performance is the difference between the value of  
271 the solution of an instance and the expected value of a knapsack filled by randomly selecting  
272 items, normalised by the latter. It is a measure of economic performance relative to a random



273 (skill-less) algorithm. The mean value of this shortfall measure was 79.7% ( $SD = 35.0\%$ ). This  
274 measure, too, varied significantly by both instance ( $min = 69.7\%$ ,  $M = 79.6\%$ ,  $max = 89.4\%$ ,  $SD$   
275  $= 6.2\%$ ) and participant ( $min = 36.4\%$ ,  $M = 79.8\%$ ,  $max = 94.3\%$ ,  $SD = 13.5\%$ ). The fact that  
276 this measure is significantly above 0 (one-sample t-test,  $t_{(307)} = 39.893$ ,  $P < 0.001$ ) is another  
277 indication that human participants performed better than a skill-less (random) algorithm.

## 278 **2.3 Effort and performance**

279 Next, we examined the relation between effort and performance in more detail. One measure of  
280 effort spent on an instance is the number of additions of items to and removals from the knap-  
281 sack, which we refer to as the length of the search path in the graph induced by the instance (SI  
282 Materials and Methods 1.2). This number can be considered as a proxy of the number of com-  
283 putations performed by the participant during an attempt, that is, a measure of computational  
284 time (analogous to *CPU time* in computing). There was a positive relation between mean path  
285 length and mean value across instances ( $r = 0.792$ ,  $P < 0.05$ ; Fig. 3A). The effect was not sig-  
286 nificant at the level of individual attempts, however (linear mixed model (LMM) with instance  
287 and participant random effects on intercept and main effect of path length,  $P > 0.05$ ; Tab. S3  
288 Model 1). There was no relation between success rate and the length of the path on average  
289 ( $r = -0.394$ ,  $P > 0.05$ ), even at the level of individual attempts (generalised linear mixed  
290 model (GLMM) with instance and participant random effects on intercept and main effect of  
291 path length,  $P > 0.05$ ; Tab. S3 Model 3).

292 Another measure of effort spent on an instance is clock time. There was no relation be-  
293 tween between mean clock time spent on an attempt and mean economic value across instances  
294 ( $r = 0.498$ ,  $P > 0.05$ ; Fig. 3B). However, there was a positive relation between time spent

295 on an attempt and the value of the attempt at the level of individual attempts (LMM with in-  
296 stance and participant random effects on intercept and main effect of clock time,  $P < 0.05$ ;  
297 Tab. S3 Model 2). Participants who spent more time on an instance achieved higher values.  
298 There was no relation between clock time spent on an instance and success rate at the level of  
299 instances ( $r = -0.446$ ,  $P > 0.05$ ), nor at the level of individual attempts (GLMM with instance  
300 and participant random effects on intercept and main effect of clock time,  $P > 0.05$ ; Tab. S3  
301 Model 4).

302 These results suggest that participants may have allocated resources (clock time and com-  
303 putational time on task) according to value. We investigated this notion in more detail. Homo  
304 economicus would be expected to keep spending effort on an attempt while marginal gain from  
305 effort is larger than marginal cost of effort. Thus, we would expect participants to keep work-  
306 ing on an attempt as long as the marginal gain per unit of time is larger than the cost of effort  
307 (which we assumed to be positive and constant). To investigate whether this was the case, we  
308 computed marginal gain from effort per unit of clock time for each attempt and averaged across  
309 all attempts. We found that that mean marginal gain per unit of clock time dropped to zero at  
310 about 60 s and remained at zero for the remainder of time on task (Fig. 3D; Fig. S2A). Given  
311 that the mean time on task was 172.0 s, as a group participants spent more than two thirds of  
312 their time on attempts at zero marginal gain. Indeed, if we assume that marginal cost of effort  
313 was strictly positive, as a group participants spent most of the time on task at a marginal net  
314 loss. The same pattern emerges when considering computational time instead of clock time  
315 (Fig. S2C).

316 We also examined how the quality of an attempt improved in item space. To this end,  
317 we computed the differences in distances  $\ell_G$  to the solution between subsequent vertices in the

318 path, which is equal to the gain in distance  $\ell_G$  between two vertices, and examined the time  
319 course of gains. The mean gain reached zero after about seven steps (Fig. S2B) or about 70 s  
320 (Fig. S2D). This means that on average, the gains in quality of attempts were achieved in the  
321 first few steps of an attempt, after which the average gain was zero. We conclude that the gains  
322 in quality in attempts in both item and value space appeared in the first third to quarter of an  
323 attempt, after which gains in quality remained around zero on average.

324 In summary, more time spent on an attempt was associated with a higher values achieved  
325 in the attempt, but it was not associated with a higher likelihood of success. We now turn to the  
326 question of what determined success in the instances, that is, computational performance.

## 327 **2.4 Computational performance vs. economic performance**

328 In the next step, we examined the relation between computational performance and economic  
329 performance. To this end, we compared success rates and economic values of attempts across  
330 instances. *Homo economicus* exerts effort until the marginal gain from effort is equal to the  
331 marginal cost of effort. The mean success rate can be interpreted as an index of difficulty of  
332 an instance. Assuming that marginal cost of effort is strictly positive and constant, we would  
333 expect a positive relation between success rate (computational performance) and value (eco-  
334 nomic performance) on average. That is, we would expect participants to make more money in  
335 instances with higher success rates (easy instances). However, we found the opposite to be the  
336 case: The mean value of attempts of an instance was negatively correlated with the mean suc-  
337 cess rate for the instance, that is, participants generated less value in easy instances compared  
338 to difficult instances (Pearson correlation  $r = -0.838$ ,  $P < 0.01$ ; Fig. 3C). This means that  
339 participants on average made more money on difficult instances. Note that for a given instance,

340 correct attempts will always be worth more than incorrect attempts. The same applies for a  
341 given participant.

## 342 **2.5 Variation in computational performance**

343 We found significant variation in success rates (computational performance) across instances  
344 and also that success did not vary with time spent on those instances (SI Results 2.2). We  
345 then investigated whether success in instances was related to instance properties, in particular  
346 various measures of their computational complexity and graph topology.

347 We first examined the relation between success and various measures of the size of the  
348 instances. Success rates were negatively related to the number of vertices in the instance graph  
349 at the level of attempts (GLMM with participant random effects on intercept and main effect of  
350 number of vertices,  $P < 0.001$ , Tab. S4 Model 1) but not at the level of instances ( $r = -0.333$ ,  
351  $P > 0.05$ ). They were also negatively correlated with the number of terminal nodes at the level  
352 of individual attempts (GLMM with participant random effects on intercept with main effect  
353 of number of terminal vertices,  $P < 0.01$ , Tab. S4 Model 2) but not at the level of instances  
354 ( $r = -0.310$ ,  $P > 0.05$ ).

355 Next, we examined the relation between computational performance and computational  
356 complexity of the instance. Mean success rate was not related to input size ( $I \log_2 C$ ) of in-  
357 stances ( $r = -0.184$ ,  $P > 0.05$ ), even at the level of individual attempts (GLMM with partic-  
358 ipant random effects on intercept and main effect of input size,  $P > 0.05$ ; Tab. S4 Model 3).  
359 However, we found that success rates were negatively related to Sahni- $k$  both at the level of  
360 instances and at the level of individual attempts ( $r = -0.689$ ,  $P = 0.059$ ; Fig. 4A; GLMM  
361 with participant random effects on intercept and main effect of Sahni- $k$ ,  $P < 0.001$ ; Tab. S4

362 Model 4). The success rate of the instance with  $k = 0$ , that is, the instance that could be solved  
363 with the greedy algorithm, was 74.4% whereas the success rate for the instance with the highest  
364  $k$  ( $k = 4$ ) was 2.7%. This suggests that there was a negative relation between computational  
365 complexity of the instances and success rate. We also found a negative relation between be-  
366 tween success rates and the Pearson correlation of item values and weights, but only at the level  
367 of individual attempts when controlling for participant random effects, and not at the level of  
368 instances ( $r = -0.637$ ,  $P = 0.089$ ; GLMM with participant random effects on intercept and  
369 main effect of Pearson correlation,  $P < 0.05$ ; Tab. S4 Model 5) but the value-weight correla-  
370 tion could not explain variation in performance that was not captured by Sahni- $k$  (GLMM with  
371 participant random effects on intercept, interaction Sahni- $k \times$  Pearson correlation,  $P > 0.05$ ;  
372 Tab. S4 Model 6). These results suggest that computational performance in the instances was  
373 strongly related to the size of the search problem induced by the instance (size of the search  
374 space) as well as computational complexity of the instance. They provide indications of what  
375 search strategies or algorithms participants may have used and where their searches for solutions  
376 broke down.

## 377 **2.6 How did participants search?**

378 To examine participants' search strategies in more detail, we considered the search paths dur-  
379 ing individual attempts, that is, the sequence of additions of items to and removal from the  
380 knapsack. From this sequence we can reconstruct the state of the knapsack at any point in time,  
381 which can be mapped on the instance graph as a search path. The average number of steps (item  
382 additions/removals) in participants' search paths was 33.3 ( $SD = 22.1$ ). During their search,  
383 participants visited 4.0 terminal vertices (maximally admissible knapsacks) on average.

384 First, we computed the proportion of vertices and terminal vertices in the graph induced by  
385 an instance that participants visited during their search. The mean proportion of unique vertices  
386 visited by participants was 3.6%, with significant variation across instances ( $min = 0.6%$ ,  $max =$   
387  $7.6%$ ,  $SD = 2.6%$ ; Fig. S4A). As a group, they visited 42.1% of vertices of the instance graph on  
388 average ( $min = 10.2%$ ,  $max = 74.8%$ ,  $SD = 24.5%$ ; Fig. S4B). This means that while individual  
389 participants only visited a very small proportion of the graph, as a group they visited a large part  
390 of it. This suggests that there was significant heterogeneity in search strategies. In addition, in  
391 all but one instance, at least one participant found the solution, which means that as a group,  
392 participants searched successfully whereas individually they did not. The mean proportion of  
393 vertices visited by participants was negatively correlated with the total number of vertices in  
394 the graph ( $r = -0.888$ ,  $P < 0.01$ ) and so was the proportion of vertices visited by the group  
395 ( $r = -0.870$ ,  $P < 0.01$ ).

396 We found a similar pattern for the proportion of unique terminal vertices visited by partic-  
397 ipants. The mean proportion of terminal vertices visited by participants was 4.6%, with signifi-  
398 cant variation across instances ( $min = 0.8%$ ,  $max = 10.6%$ ,  $SD = 3.1%$ ; Fig. S4C). As a group,  
399 they visited 52.1% of terminal vertices on average ( $min = 12.5%$ ,  $max = 74.0%$ ,  $SD = 20.7%$ ;  
400 Fig. S4D). There was also a large degree of heterogeneity in the number of terminal vertices  
401 submitted at the end of an attempt. The mean number of unique terminal vertices submitted by  
402 participants was 13.9 ( $min = 7$ ,  $max = 30$ ,  $SD = 7.4$ ). The mean proportion of terminal vertices  
403 visited by participants was negatively correlated with the total number of terminal vertices in  
404 the graph ( $r = -0.861$ ,  $P < 0.01$ ) and so was the proportion of vertices visited by the group  
405 ( $r = -0.948$ ,  $P < 0.001$ ).

406 We conclude that while individual participants only explored a relatively small part of the

407 search space, as a group they explored a large part of it. However, at the level of instances, we  
408 found that there was no relation between success in an attempt and the proportion of vertices  
409 participants visited as a group ( $r = 0.388$ ,  $P > 0.05$ ) or the proportion of terminal vertices  
410 visited ( $r = 0.230$ ,  $P > 0.05$ ). The same was the case at the level of participants within  
411 instances (GLMM with participant and instance random effects on intercept and main effect  
412 of proportion of (terminal) vertices visited,  $P > 0.05$ ; Tab. S5 Models 1 and 2). That is,  
413 participants who explored a larger part of the search space in an instance were not more likely  
414 to be successful in the instance. We conclude that the *extent* of search was not associated with  
415 computational performance.

416 We also investigated the relation between the extent of search and value. There was no  
417 relation between value in an attempt and the proportion of vertices participants visited as a  
418 group ( $r = -0.596$ ,  $P > 0.05$ ) or the proportion of terminal nodes visited ( $r = -0.048$ ,  
419  $P > 0.05$ ) at the level of instances nor at the level at individual attempts (LMM with participant  
420 and instance random effects on intercept and main effect of proportion of (terminal) vertices  
421 visited,  $P > 0.05$ ; Tab. S5 Models 3 and 4).

422 Next, we examined the *quality* of search. To do so, we compared the quality of the vertices  
423 visited to the average quality of the vertices in the graph. If participants picked vertices at  
424 random, then the quality of the vertices visited would be equal to the average quality of all  
425 vertices in the graph. First, we looked at the distance to the solution  $\ell_G$  of vertices visited (SI  
426 Materials and Methods 1.2). For each attempt, we computed  $\ell_G$  of each of the vertices visited  
427 and computed the mean of those values. This gives us the mean of  $\ell_G$  of all vertices visited.  
428 From it we subtracted the mean of  $\ell_G$  of *all* vertices in the graph induced by the instance. The

429 mean value of this difference was -1.230, which was significantly below zero (one-sample t-test,  
430  $t_{(307)} = -17.461, P < 0.001$ ). It implies that the quality of vertices visited by participants was  
431 significantly better than the average quality of vertices in the instances.

432 We found that the gains in quality of an attempt occurred mainly in the first stage of an  
433 attempt (SI Results 2.4). To examine in more detail the notion that only the earlier but not  
434 the later stages of the search were beneficial, we considered the terminal vertices visited by  
435 participants during their attempts. More specifically, we compared the quality in item space  
436 of the first terminal vertex visited to the quality of the last terminal vertex visited. The first  
437 terminal vertex is the first full knapsack (set of items) a participant assembled and the last  
438 terminal vertex is the knapsack submitted. We measured quality of a vertex  $i$  by its distance  
439 to the solution vertex  $s$ ,  $\ell_G(i, s)$  (SI Materials and Methods 1.2). The mean distance of the  
440 first terminal vertex visited to the solution vertex across instances was 3.699 ( $min = 2.526, max$   
441  $= 5.350, SD = 0.876$ ). In comparison, the mean number of items in the solution was 5.500  
442 ( $min = 3, max = 9, SD = 1.871$ ). The mean distance of the last terminal vertex was 2.628  
443 ( $min = 0.763, max = 4.800, SD = 1.163$ ). This means that there was a greater improvement in  
444 quality between initial vertex and first terminal vertex than between first and last terminal vertex  
445 visited. In addition, the mean difference between the terminal vertices visited, that is  $\ell_G(i, j)$   
446 where  $i$  and  $j$  are two subsequent terminal vertices on the search path, was 1.809 ( $min = 1.593,$   
447  $max = 2.052, SD = 0.131$ ). Note that participants visited about 4 terminal vertices on average.  
448 This means that the mean distance between the terminal vertices visited was higher than the  
449 reduction in distances to the solution between first terminal vertex to last terminal vertex. It  
450 suggests that many of the changes in the sets of items between first and last terminal vertex did  
451 not result in a reduction of the distance to the solution.



452 We also computed the proportion of participants that had visited the solution vertex for  
453 each step in the search path. This gives us, for each step in the search, the proportion of partic-  
454 ipants who visited the solution vertex by that step. The mean proportion of participants across  
455 instances who visited the solution vertex was 39.6% ( $min = 2.7\%$ ,  $max = 79.5\%$ ,  $SD = 20.6\%$ ),  
456 which is slightly higher than the mean success rate. The mean number of steps across instances  
457 until the first participant visited the solution vertex was 7.2 ( $min = 4$ ,  $max = 12$ ,  $SD = 2.9$ ;  
458 Fig. S5), compared to a mean number of steps in the search path of 33.3. Across instances,  
459 among all participants who visited the solution vertex, the mean number of steps to the first  
460 visit was 18.0 ( $min = 8.6$ ,  $max = 35.9$ ,  $SD = 8.1$ ). This means that among those participants  
461 who visited the solution vertex, the first participant to visit tended to be substantially faster than  
462 the average, another sign of heterogeneity in search strategies. However, most participants who  
463 visited the solution vertex kept searching before they submitted their solution. The mean num-  
464 ber of steps between the first visit of the solution vertex and submission of the attempt was 21.1  
465 ( $min = 6.1$ ,  $max = 26.2$ ,  $SD = 17.0$ ). In that time, many participants visited the solution vertex  
466 multiple times before they submitted an attempt. The mean number of visits across instances,  
467 among those participants who visited the solution vertex at least once, was 2.8 ( $min = 2.0$ ,  $max =$   
468  $5.1$ ,  $SD = 0.9$ ). In addition, in some of the instances the proportion of participants who visited  
469 the solution vertex was higher than the success rate in the instances (Fig. S5). This suggests  
470 that some of the participants visited the solution vertex but submitted another set of items in the  
471 attempt, a point probably related to the NP completeness property of knapsack problems, which  
472 we examine in more detail below (SI Results2.9).

## 473 **2.7 Which search algorithms did participants use?**

474 Performance data in combination with information about the search path allows certain in-  
475 ferences about the type of search algorithm participants may have used. The relatively low  
476 computational performance together with the short average length of the search path and small  
477 fraction of the instance graphs participants explored, suggests that participants did not use any  
478 of the uninformed, exhaustive search algorithms.

479 On the other hand, participants' computational performance was substantially higher than  
480 that of a random algorithm, suggesting that participants used an informed (rule- or heuristic-  
481 based) algorithm. The fact that performance was well below 100%, rules out dynamic pro-  
482 gramming. This conclusion is further supported by the absence of a relation between success  
483 rate and input size of the problem (SI Results 2.5). It is more likely that participants used some  
484 sort of approximation algorithm. The finding that success rates decreased with the Sahni- $k$  of  
485 instances suggests that participants were using an algorithm of low computational complexity  
486 (SI Results 2.5). The relatively high success rate in instances with a Sahni- $k$  of 0 suggests  
487 that the algorithm used was similar to the greedy algorithm. To investigate this possibility in  
488 more detail, we examined the sequence of additions of items to and removals of items from  
489 the knapsack. For each instance, we ordered the items in the various instances in decreasing  
490 order of value-to-weight ratio and computed the frequencies with which the items at each rank  
491 were chosen in the various steps of participants' sequences. If all participants used the greedy  
492 algorithm, then the items with the highest value-to-weight ratio of each instance would have  
493 been chosen in the first step, the item with the second highest value-to-weight ratio would have  
494 been chosen in the second step, and so on.

495 We found that across all participants and all instances, the items with the highest value-  
496 to-weight ratios were chosen most often in the first few steps. For example, the three items  
497 with the highest value-to-weight ratio were chosen in 15.6% of cases in the first three steps on  
498 average, while the mean frequency for the next seven items was 6.8% (Fig. 3B). In addition, the  
499 frequencies with which the items were chosen decreased with the number of steps. The three  
500 items with the highest value-to-weight ratios were chosen in 9.4% of cases in steps four to 10  
501 on average, compared to 15.6% in the first three steps (Fig. 4D). These patterns were similar  
502 across all instances (Fig. S3). They suggest that participants selected the items with the highest  
503 value-to-weight ratios first when filling the knapsack, similar to the greedy algorithm. However,  
504 there was considerable variation in the order with which items were chosen, which suggests  
505 that participants either did not follow the greedy algorithm exactly or that not all participants  
506 followed the greedy algorithm.

507 Several other findings provide further support for the claim that participants did not use  
508 the greedy algorithm. Firstly, participants' performance was substantially higher than that of  
509 the greedy algorithm (it would only have found the solution in one of the instances whereas  
510 participants solved 37.4% of instances on average). Secondly, the greedy algorithm fills the  
511 knapsack by selecting items into the knapsack in decreasing order of their value-to-weight ratio,  
512 until the knapsack is full. This means that the greedy algorithm would have terminated attempts  
513 after 6.6 steps on average. The average number of steps in participants' sequences (searches)  
514 was 33.3, however ( $SD = 22.1$ ).

515 These results suggest that participants were more likely to have used an algorithm similar  
516 to branch-and-bound that starts the search by filling the knapsack with the greedy algorithm  
517 and then searches for improvements by systematically removing and adding items in search

518 for higher value knapsacks. Since Sahni- $k$  is a measure of deviation of a solution from the  
519 greedy algorithm, we would expect that participants who tried to replace multiple items in the  
520 first full knapsack to be more successful, at least for instances where this was needed, that is,  
521 for high Sahni- $k$  instances. Therefore, we tested whether computational performance could be  
522 explained, not only by Sahni- $k$ , but also by the interaction between Sahni- $k$  and the number of  
523 items participants replaced on average after reaching the first full knapsack. We measured the  
524 latter as the length of the shortest path between two full knapsack attempts, that is, between two  
525 subsequent terminal vertices (SI Materials and Methods 1.2). The interaction term was indeed  
526 significant (GLMM with random effect for instances on intercept and interaction of Sahni- $k \times$   
527 mean distance,  $P < 0.01$ ; Table S6). However, the fact that the values of the knapsacks did  
528 not increase over time for the last two thirds of participants' searches (Fig. S2A, C) suggests  
529 that participants did not use the branch-and-bound algorithm, at least not in its exact form. The  
530 high average length of sequences and the low average number of terminal vertices visited also  
531 suggests that participants did not use the Sahni algorithm.

## 532 **2.8 Path dependence in search**

533 Next, we investigated whether there was path dependence in the search paths. To this end, we  
534 examined the sequence of terminal vertices (maximally admissible knapsacks) visited by par-  
535 ticipants during an attempt. First, we tested whether the distance  $\ell_G(i, s)$  of the first terminal  
536 vertex  $i$  to the solution  $s$  was predictive of the distance of the last terminal vertex. We esti-  
537 mated a LMM with distance of the last terminal vertex as dependent variable and distance of  
538 the first terminal vertex as independent variable, with random effects on intercept for instances.  
539 We found that the distance of the first terminal vertex was predictive of the distance of the last

540 terminal vertex, and hence of success (LMM with instance and participant random effects on  
541 intercept and main effect of distance of first terminal vertex,  $P < 0.001$ ; Tabel S7 Model 1).  
542 This suggests that quality of an attempt (distance of the last terminal vertex to the solution)  
543 was path-dependent. We also found that the change in distances between first and last terminal  
544 vertex was predictive of the distance of the last terminal node (LMM with instance and partic-  
545 ipant random effects on intercept and main effect of change in distance,  $P < 0.001$ ; Table S7  
546 Model 2), which means that the quality of the search increased the likelihood of success.

547 We examined the notion of path dependence further by investigating to what extent there  
548 was a tendency not to eliminate incorrect items that were added early on, and whether this  
549 determined computational performance. To this end, we considered the distribution of the age  
550 of incorrect items that were eventually deleted (Fig. 5). We defined age as a fraction of number  
551 of steps taken since the beginning of an attempt (age equals 1 if the item was the first added to  
552 the knapsack). Most deleted incorrect items were added very recently ( $M = 0.2920$ ,  $SEM =$   
553  $0.0001$ ); only rarely did participants eliminate incorrect items that were added to the knapsack  
554 early on. A similar pattern emerged for correct items that were deleted ( $M = 0.2352$ ,  $SEM =$   
555  $0.0001$ ; Fig. 5). Mean age of correct items was significantly higher than age of incorrect items  
556 (two-sample  $t$ -test,  $t = 6.98$ ,  $P < 0.001$ ) and their distributions were significantly different  
557 (Kolmogorov-Smirnov test for independence of samples,  $D = 0.10$ ,  $P < 0.001$ ).

558 Further evidence of path dependence in the search is provided by an analysis of the indi-  
559 vidual items in the first and last terminal vertices, respectively. We computed the probability  
560 of a correct item in the first terminal vertex being left in the knapsack, conditional on it being  
561 correct. In the absence of path dependence, this conditional probability would be 0.5, that is,  
562 the probability of a correct item being left in the knapsack would be at chance. The mean of this

563 probability across instances was 0.871 ( $min = 0.734$ ,  $max = 0.945$ ,  $SD = 0.063$ ). On the other  
564 hand, the probability of an incorrect item in the first terminal vertex being removed from the  
565 knapsack, conditional on it being incorrect, was 0.518 on average ( $min = 0.364$ ,  $max = 0.833$ ,  
566  $SD = 0.136$ ). These results further support the notion that there was path dependence in the  
567 search paths but that it was largely due to correct items in the first terminal vertex being left in  
568 the knapsack, as opposed to incorrect items being left in the knapsack. The fact that the proba-  
569 bility of incorrect items in the first terminal vertex being removed during subsequent stages of  
570 the search was at chance, suggests that participants were not very good at identifying incorrect  
571 items.

## 572 **2.9 Did participants solve the decision problem?**

573 In the theory of computation, a distinction is made between computational problems and de-  
574 cision problems. A computational problem is a problem that can be solved by mechanical  
575 application of mathematical operations, such as an algorithm. A decision problem is a special  
576 case of a computational problem whose answer is either yes or no. In the case of the knapsack  
577 problem, the decision problem is 'Can a value of at least  $V$  be achieved without exceeding the  
578 weight  $C$ ?' Put differently, the computational problem is the problem of finding the optimal  
579 solution of an instance (also referred to as *optimisation problem*), whereas the decision prob-  
580 lem is the problem of verifying that a candidate solution is the actual solution of an instance.  
581 The decision problem form of the knapsack problem is NP-complete, that is, there is no known  
582 polynomial algorithm which can verify that the decision is true (2). Given that the decision  
583 problem is NP-complete, the optimisation problem of the knapsack problem is NP-hard, that is,  
584 there is no polynomial algorithm for solving the optimisation problem (17).

585        So far, we have analysed the computational problem of the knapsack problem. We now  
586 examine whether those participants who solved the optimisation problem were also able to solve  
587 the decision problem, that is, whether they were able to verify that a candidate solution is the  
588 solution of the problem. As reported in the previous section, those participants who visited the  
589 solution vertex at least once tended to visit it several times. This suggests that those participants  
590 did not know that they had found the solution, that is, they could not solve the decision problem.  
591 A participant who was able to solve the decision problem would have known that the solution  
592 vertex is indeed the highest value vertex, and hence would have submitted this set of items in  
593 their attempt. Moreover, considering only those attempts in which the participant visited the  
594 solution vertex at least once, in 6.5% of cases the participant subsequently submitted another  
595 set of items (of inferior value). These participants definitely did not solve the decision problem.

596        We also examined how participants performed on subsequent attempts of the same in-  
597 stance. Every participant attempted the same instance twice, with one attempt immediately  
598 following the other. A participant who solved the decision problem would be expected to re-  
599 member the solution and therefore also solve the second attempt. Thus, we computed the num-  
600 ber of times participants solved the first attempt of the same instance or the second or both. The  
601 percentage of participant  $\times$  instance pairs in which participants found the solution in at least  
602 one attempt was 51.0%. In 22.9% of cases, participants solved both the first and the second  
603 attempt. In 17.6% of cases, they only solved the second attempt, and in 10.4% of cases they  
604 only solved the first attempt. Success in first and second attempt was not independent ( $\chi^2$  test,  
605  $\chi^2_{(2,153)} = 23.3, P < 0.001$ ). We would expect the number of successful second attempts to be  
606 higher than the number of successful first attempts, as participants had already explored part of  
607 the search space. However, we would not expect there to be any cases in which a participant

608 solved an instance in the first attempt but not in the second attempt. The fact that of all partic-  
609 ipants who solved instances at least once, 20.5% only solved the instance in the first attempt  
610 but not in the second attempt, which indicates that those participants did not solve the decision  
611 problem, that is, they did not know that they had found the solution.

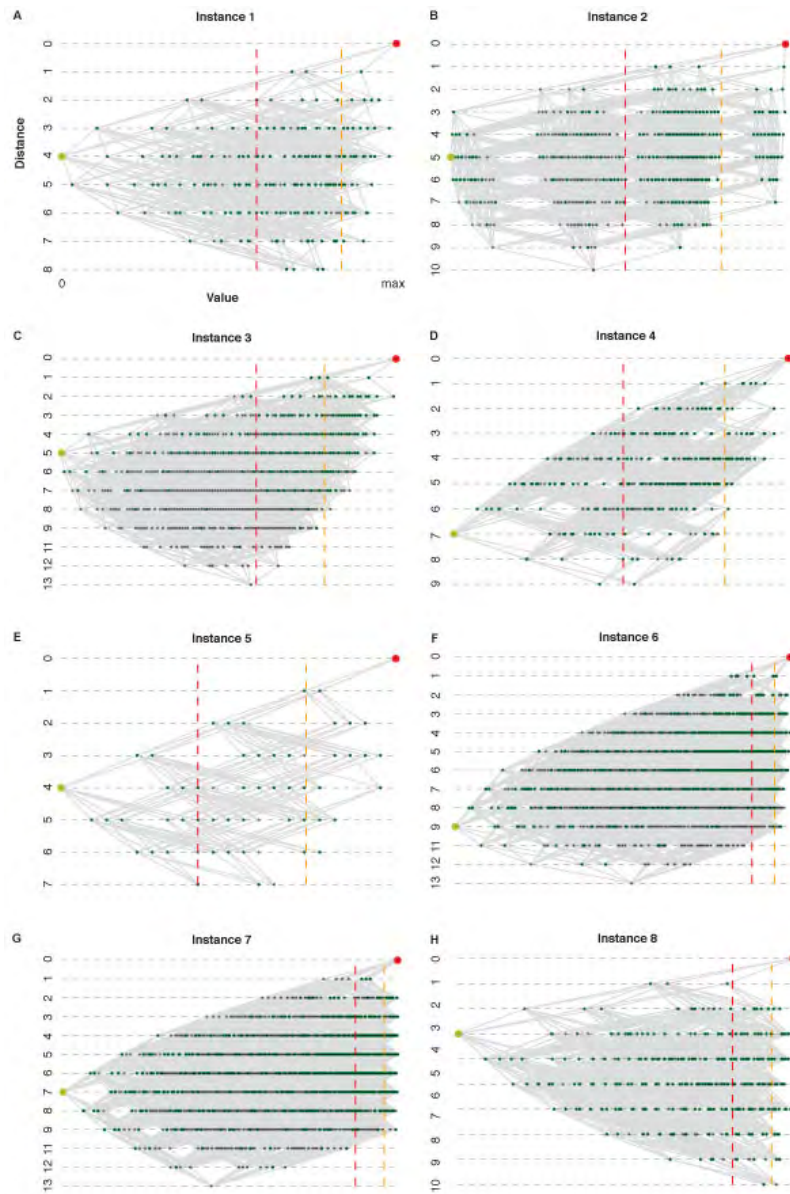


## 612 **References**

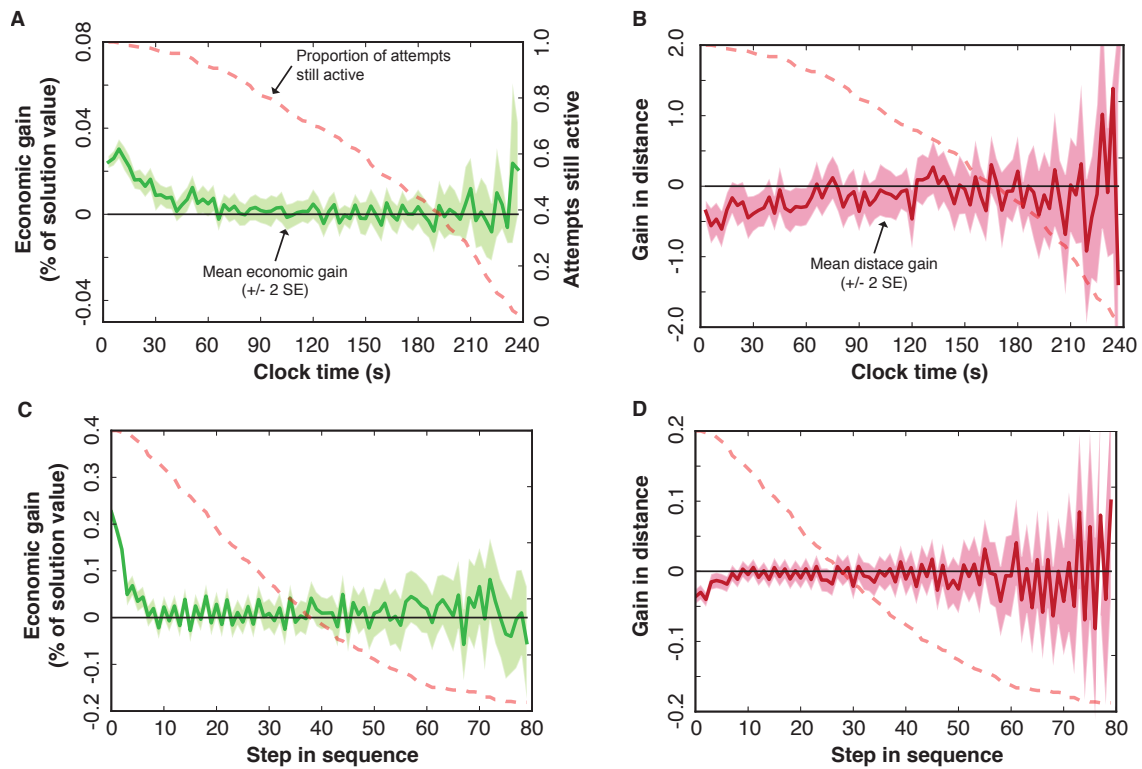
- 613 [1] Mathews GB (1897) On the partition of numbers. *Proceedings of the London Mathemati-*  
614 *cal Society* 28:486–490.
- 615 [2] Kellerer H, Pferschy U, Pisinger D (2004) *Knapsack Problems*. (Springer Science &  
616 Business Media).
- 617 [3] Cook SA (1983) An overview of computational complexity. *Communications of the ACM*  
618 26(6):400–408.
- 619 [4] Karp RM (1972) Reducibility among Combinatorial Problems in *Complexity of Computer*  
620 *Computations*. (Springer US), pp. 85–103.
- 621 [5] Fortnow L (2009) The status of the P versus NP problem. *Communications of the ACM*  
622 52(9):78–86.
- 623 [6] Diestel R (2006) *Graph Theory*. (Springer).
- 624 [7] Navarro G (2001) A guided tour to approximate string matching. *ACM Computing Surveys*  
625 33(1):31–88.
- 626 [8] Dasgupta S, Papadimitriou C, Vazirani U (2006) *Algorithms*. (McGraw-Hill, New York,  
627 NY).
- 628 [9] Cormen TH, Leiserson C, Rivest RL, Stein C (2001) *Introduction to Algorithms*. (MIT  
629 Press, Cambridge, MA).
- 630 [10] Marr D (1982) *Vision*. (W.H. Freeman, San Francisco).

- 631 [11] Russell S, Norvig P (2014) *Artificial Intelligence*. (Pearson, Harlow).
- 632 [12] Newell A, Simon H (1972) *Human Problem Solving*. (Prentice-Hall, Englewood Cliffs).
- 633 [13] Land AH, Doig A (1960) An automatic method for solving discrete optimization prob-  
634 lems. *Econometrica* 28(3):497–520.
- 635 [14] Sahni S (1975) Approximate algorithms for the 0-1 knapsack problem. *Journal of the*  
636 *ACM* 22:115–124.
- 637 [15] Meloso D, Copic J, Bossaerts P (2009) Promoting intellectual discovery: patents versus  
638 markets. *Science* 323(5919):1335–1339.
- 639 [16] Pisinger D (2004) Where are the hard knapsack problems? *Computers & Operations*  
640 *Research* 32:2271–2284.
- 641 [17] Moore C, Mertens S (2011) *The Nature of Computation*. (Oxford University Press, Ox-  
642 ford).

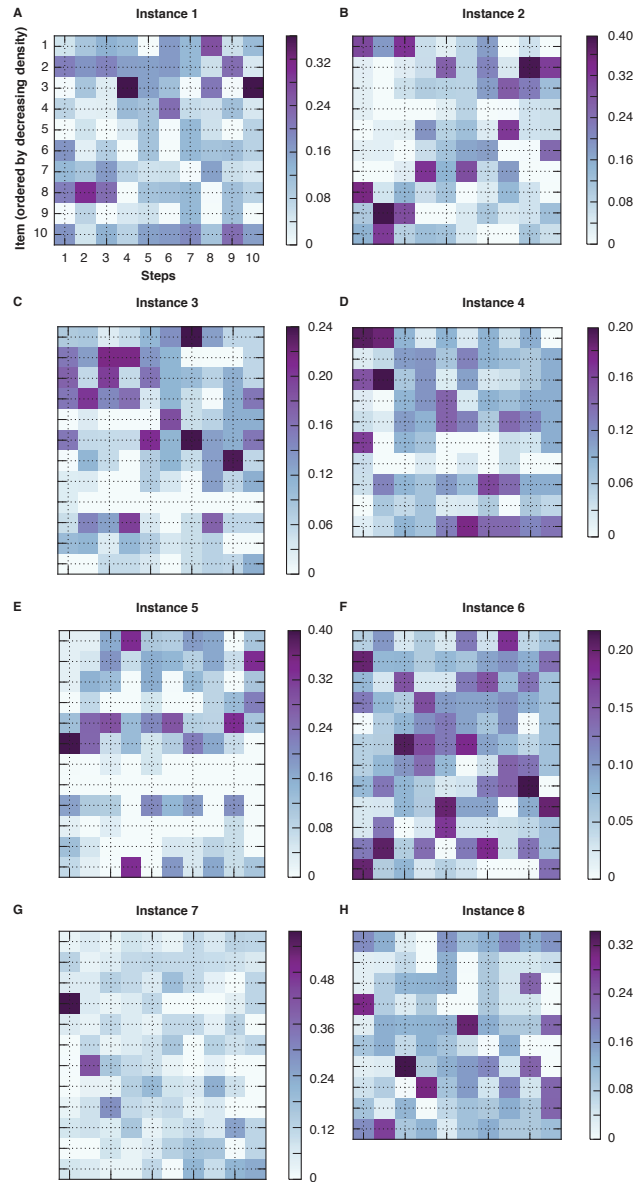
643 **3 Figures**



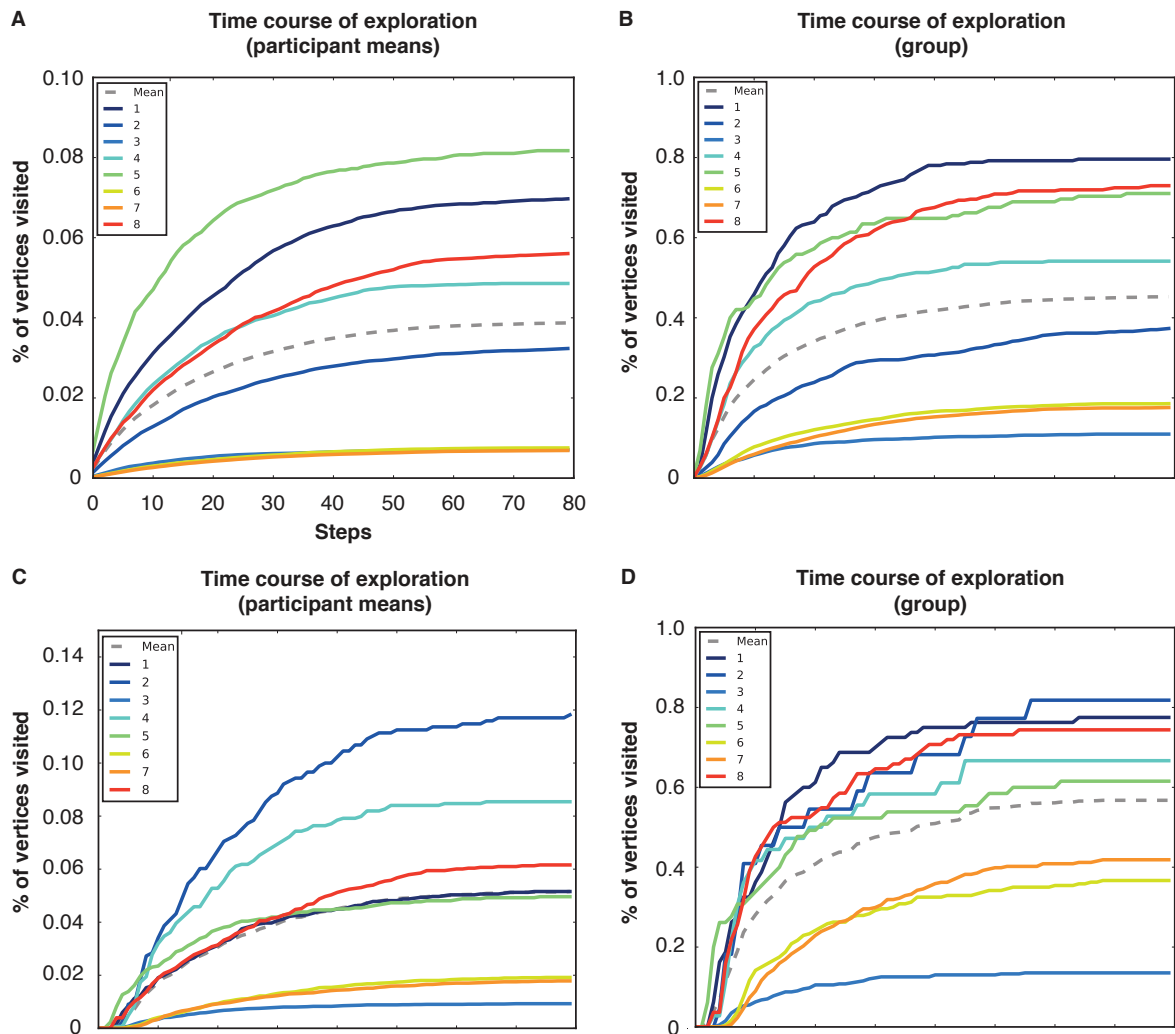
**Figure S1.** Instance graphs for each instance in the task. (A–H) Graph induced by the instance (SI Materials and Methods 1.2). Each vertex represents an admissible set of items. The initial vertex (empty set) is coloured in yellow and the solution vertex is coloured in red. Two vertices are connected by an edge if one vertex can be reached from the other by adding or removing one item. The position of a vertex on the abscissa is determined by the total value of the set of items represented by the vertex. The position of a vertex on the ordinate is determined by the distance  $\ell_G$  (shortest path length) of the vertex to the solution vertex. The red dashed line indicates the lowest value of any terminal vertex and the yellow dashed line indicates the mean value of all terminal vertices. The vertices by participants during their attempts are coloured in green. The set of available items in each instance is provided in Table S1 and some key properties of the instance graphs are provided in Table S2.



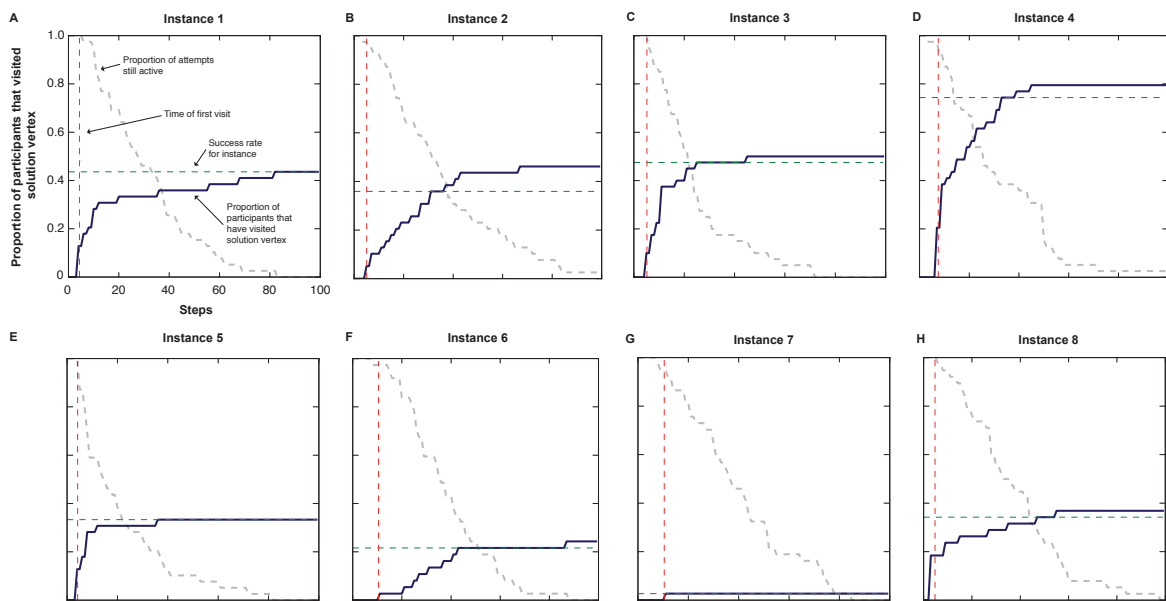
**Figure S2.** Time courses of value gain and distance gain. (A) Time course of mean value gain per unit of clock time. The mean was computed over all attempts. (B) Time course of distance gain per unit of clock time. The plot shows that mean change in distances  $\ell_G$  per unit of clock time (SI Materials and Methods 1.2). (C) Time course of mean value gain per sequence step. (D) Time course of distance gain per sequence step.



**Figure S3. Time courses of choice frequencies for individual items.** *A–H*, The items available in an instance were sorted in reverse order of their density (value-to-weight ratio). The heat map shows choice frequencies for the items for the first 11 steps in the search path (SI Materials and Methods 1.2). If the greedy algorithm was used, off-diagonal entries would be zero.



**Figure S4.** Time courses of exploration. (A) Proportion of vertices visited by individual participants in each of the instances. The lines represent the mean of participant values in each of the instance. (B) Proportion of vertices visited by all participants. The lines represent the proportion of vertices represented by the set of vertices visited by all participants at a particular step. (C) Proportion of terminal vertices visited by individual participants in each of the instances. (D) Proportion of terminal vertices visited by all participants.



**Figure S5.** Time courses of solution vertex visits. (A–H) The solid blue line shows the proportion of participants who have visited the solution vertex at a particular step during the attempt. The red dashed line indicates the step number at which the first participant visited the solution vertex. The green dashed line indicates the proportion of participants whose attempt was correct, that is, whose final set of items selected was identical to the set of items in the solution.



## 4 Tables

**Table S1.** Available items and capacity (maximum weight) for each of the instances used in the experiment. Density is defined as the ratio of value to weight.

<b>Instance 1</b>	<b>Items</b>											
<i>Capacity: 1,900</i>	1	2	3	4	5	6	7	8	9	10		
Value	500	350	505	505	640	435	465	50	220	170		
Weight	750	406	564	595	803	489	641	177	330	252		
Density	0.67	0.86	0.90	0.85	0.80	0.89	0.73	0.28	0.67	0.67		
<hr/>												
<b>Instance 2</b>	<b>Items</b>											
<i>Capacity: 1,044</i>	1	2	3	4	5	6	7	8	9	10		
Value	300	350	400	450	47	20	8	70	5	5		
Weight	205	252	352	447	114	50	28	251	19	20		
Density	1.46	1.39	1.14	1.01	0.41	0.40	0.29	0.28	0.26	0.25		
<hr/>												
<b>Instance 3</b>	<b>Items</b>											
<i>Capacity: 850</i>	1	2	3	4	5	6	7	8	9	10	11	12
Value	15	14	3	3	10	9	28	28	31	25	24	1
Weight	129	144	77	77	66	60	184	184	229	184	219	72
Density	0.12	0.10	0.04	0.04	0.15	0.15	0.15	0.15	0.14	0.14	0.11	0.01
<hr/>												
<b>Instance 4</b>	<b>Items</b>											
<i>Capacity: 1,500</i>	1	2	3	4	5	6	7	8	9	10		
Value	37	72	106	32	45	71	23	44	85	62		
Weight	50	820	700	46	220	530	107	180	435	360		
Density	0.74	0.09	0.15	0.70	0.20	0.13	0.21	0.24	0.20	0.17		

<b>Instance 5</b>	<b>Items</b>											
<i>Capacity: 14</i>	1	2	3	4	5	6	7	8	9	10	11	12
Value	2	3	4	5	6	9	8	7	6	5	8	9
Weight	3	4	6	3	5	13	6	9	2	4	7	7
Density	0.67	0.75	0.67	1.67	1.20	0.69	1.33	0.78	3.00	1.25	1.14	1.29
<hr/>												
<b>Instance 6</b>	<b>Items</b>											
<i>Capacity: 3,800</i>	1	2	3	4	5	6	7	8	9	10	11	12
Value	107	35	120	206	88	34	28	110	88	101	74	53
Weight	599	196	670	1204	502	202	145	600	453	601	404	299
Density	0.18	0.18	0.18	0.17	0.18	0.17	0.19	0.18	0.19	0.17	0.18	0.18
<hr/>												
<b>Instance 7</b>	<b>Items</b>											
<i>Capacity: 1,300</i>	1	2	3	4	5	6	7	8	9	10	11	12
Value	201	84	113	303	227	251	129	147	86	127	144	167
Weight	192	80	106	288	212	240	121	140	82	120	137	160
Density	1.05	1.05	1.07	1.05	1.07	1.05	1.07	1.05	1.05	1.06	1.05	1.04
<hr/>												
<b>Instance 8</b>	<b>Items</b>											
<i>Capacity: 265</i>	1	2	3	4	5	6	7	8	9	10		
Value	31	141	46	30	74	105	119	160	59	71		
Weight	21	97	32	21	52	75	86	116	43	54		
Density	1.48	1.45	1.44	1.43	1.42	1.40	1.38	1.38	1.37	1.31		

**Table S2.** Properties of the instances of the 0-1 knapsack problem used in the experiment.

<b>Instance</b>	<b>1</b>	<b>2</b>	<b>3</b>	<b>4</b>	<b>5</b>	<b>6</b>	<b>7</b>	<b>8</b>
Number of available items	10	10	12	10	12	12	12	10
Pearson correlation value/weight	0.955	0.903	0.929	0.856	0.856	0.997	1.000	0.998
Number of vertices in $G$	255	691	2,278	386	145	3,273	3,640	385
Number of terminal vertices in $G$	80	22	399	36	65	240	301	82
Number of edges in $G$	796	3,018	11,100	1,439	369	17,892	20,587	1,377
Pearson correlation value/ $\ell$ of vertices	-0.10	-0.23	-0.32	-0.41	-0.32	-0.18	-0.30	0.04
Number of items in solution	4	5	5	7	4	9	7	3
Input size ( $I \log_2 C$ )	109	100	117	105	46	143	124	80
Sahni- $k$	1	3	2	0	1	1	4	3
Success rate	0.44	0.36	0.48	0.74	0.33	0.22	0.03	0.34
Mean $\ell$ of attempt	2.0	2.7	1.6	0.8	2.3	3.3	4.9	3.4
Mean attempt value (% of solution value)	0.97	0.98	0.96	0.96	0.96	0.99	0.99	0.98
Mean length of search path	31.6	37.7	25.1	33.6	21.8	38.6	40.3	38.2
Mean % of vertices visited	0.07	0.03	0.01	0.05	0.08	0.01	0.01	0.06
% of vertices visited (group)	0.80	0.39	0.11	0.54	0.71	0.19	0.18	0.17
Mean % of terminal vertices visited	0.05	0.12	0.01	0.09	0.05	0.02	0.02	0.06
% of terminal vertices visited (group)	0.78	0.82	0.14	0.67	0.62	0.37	0.42	0.74

**Table S3.** Relation between performance and effort at the level of individual attempts. Models (1) and (2) related the value achieved in an attempt, normalised by the value of the solution, to the length of the search path (1) and clock time (2). Models (3) and (4) related success in a single attempt to the length of the search path (3) and clock time spent on the attempt (4). All models had random effects for both instances and participants on the intercept.

Value	<i>Dependent variable:</i>			
	(1)	(2)	(3)	(4)
			<i>Attempt correct</i>	
			<i>Generalized linear</i>	
			<i>mixed-effects</i>	
Length of search path	0.0003 (0.0002)		-0.002 (0.007)	
Clock time on attempt (s)		0.0002* (0.0001)		0.001 (0.003)
Constant	-0.841 (0.673)	-0.662 (0.487)	0.946*** (0.012)	0.965*** (0.008)
Observations	308	308	308	308
Log Likelihood	-184.622	-184.618	436.380	436.132
Akaike Inf. Crit.	377.244	377.236	-862.760	-862.263

*Note:* \*P<0.05; \*\*P<0.01; \*\*\*P<0.001

**Table S4.** Relation between performance and size and complexity of instances. Results of estimation of generalized linear mixed models relating success in an attempt to number of vertices in the instance graph (1), number of terminal vertices (2), input size (3), Sahni- $k$  (4), Pearson correlation between values and weights of items available in the instance (5) and factorial model with Sahni- $k$  and Pearson correlation between values and weights (6). All models had random effects for participants on the intercept.

	<i>Dependent variable:</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
Number of vertices	-0.0004*** (0.0001)					
Number of terminal vertices		-0.003** (0.001)				
Input size			-0.005 (0.004)			
Sahni- $k$				-0.514*** (0.108)		0.567 (1.817)
Pearson correlation value/weight					-2.672* (1.169)	1.183 (2.535)
Sahni- $k$ × Pearson correlation						-1.138 (1.918)
Constant	-0.017 (0.475)	-0.004 (0.207)	-0.176 (0.213)	0.343 (0.249)	1.875 (1.073)	-0.755 (2.297)
Observations	308	308	308	308	308	308
Log Likelihood	-200.497	-190.612	-197.426	-188.555	-198.591	-188.376
Akaike Inf. Crit.	406.993	387.224	400.853	383.110	403.181	386.752

*Note:* \*P<0.05; \*\*P<0.01; \*\*\*P<0.001

**Table S5.** Relation between performance and extent of search. Results of estimation of generalized linear mixed models relating success in an attempt to proportion of vertices in instance graph visited (1) and proportion of terminal vertices visited (2), and results of linear mixed model relating the value of an attempt, normalised by the value of the solution, to proportion of vertices in instance graph visited (3) and proportion of terminal vertices visited (4). All models had random effects for participants on the intercept.

<i>Dependent variable:</i>				
	Attempt correct		Value of attempt	
	(1)	(2)	(3)	(4)
	<i>generalized linear mixed-effects</i>		<i>linear mixed-effects</i>	
% of vertices visited	-1.245 (4.999)		-0.001 (0.100)	
% of terminal vertices visited		2.070 (3.105)		0.098 (0.069)
Constant	-0.673 (0.473)	-0.824 (0.446)	0.974*** (0.007)	0.969*** (0.007)
Observations	308	308	308	308
Log Likelihood	-184.618	-184.429	441.095	441.668
Akaike Inf. Crit.	377.235	376.858	-872.190	-873.335

*Note:* \*P<0.05; \*\*P<0.01; \*\*\*P<0.001

**Table S6.** Relation between computational performance and variation in search. Results of estimation of generalised linear mixed models relating success in an attempt to Sahni- $k$ , the mean distance  $\ell_G$  between subsequent terminal vertices visited during an attempt (SI Materials and Methods 1.2) and the interaction between Sahni- $k$  and mean distance between subsequent terminal vertices. The model had random effects for participants on the intercept.

	<i>Dependent variable:</i>
	Attempt correct
Sahni- $k$	−1.208*** (0.294)
Mean distance between terminal vertices	−0.415 (0.264)
Interaction Sahni- $k$ and mean distance	0.355** (0.134)
Constant	1.118* (0.532)
Observations	279
Log Likelihood	−167.351
Akaike Inf. Crit.	344.701

*Note:*

\* $P < 0.05$ ; \*\* $P < 0.01$ ; \*\*\* $P < 0.001$

**Table S7.** Path dependence in search. Results of estimation of linear mixed models relating distance  $\ell_G$  between the last vertex in an attempt and the solution vertex (SI Materials and Methods 1.2), to the distance of the first terminal vertex in an attempt from the solution vertex (1) and the mean change in distances between subsequent terminal vertices (2). The model had random effects for participants and instances on the intercept.

	<i>Dependent variable:</i>	
	$\ell_G$	
	(1)	(2)
Distance first to last terminal vertex	0.445*** (0.051)	
Change in distances		1.043*** (0.059)
Constant	0.989*** (0.363)	4.234*** (0.283)
Observations	304	304
Log Likelihood	-624.094	-550.643
Akaike Inf. Crit.	1258.189	1111.287

*Note:*

\*P<0.05; \*\*P<0.01; \*\*\*P<0.001