

NC STATE UNIVERSITY

Center for Environmental and Resource Economic Policy
College of Agriculture and Life Sciences
<https://cenrep.ncsu.edu>

Stochastic Dynamic Programming without Transition Matrices

Paul L. Fackler

Center for Environmental and Resource Economic Policy
Working Paper Series: No. 18-018
September 2018

Suggested citation: Fackler, Paul L. (2018). Stochastic dynamic programming without transition matrices. (CEnREP Working Paper No. 18-018). Raleigh, NC: Center for Environmental and Resource Economic Policy.



1 Stochastic dynamic programming without transition matrices

2 Paul L. Fackler

3 Professor, North Carolina State University

4 9/17/2018

5 **Abstract:** Discrete dynamic programming, widely used in addressing optimization over time, suffers
6 from the so-called curse of dimensionality, the exponential increase in problem size as the number of
7 system variables increases. One method to partially address this problem is to avoid the use of state
8 transition probability matrices, which grow in the square of the size of the state space. This can be done
9 through the use of expected value (EV) functions, which compute the expectation of functions of the
10 future state variables conditioned on current variables. Two ways that this leads to potential gains arise
11 when the state transition can be broken into separate phases and when the transitions for different state
12 variables are conditionally independent. Both of these situations arise in models that are used in natural
13 resource management and are illustrated with several examples.

14 **Keywords:** dynamic programming, computational efficiency, SPOMs, harvest management, Kronecker
15 products

16

17 Discrete dynamic programming (DDP) is a fundamental tool for making good decisions
18 concerning dynamically changing systems. For a gentle introduction see Marescot, et al. (2013) and for
19 more in-depth discussions see Puterman (1994) or Rust (2008). A significant limitation of DDP, the so-
20 called curse of dimensionality, arises due to the exponential increase in the problem size as the number of
21 variables increases (Powell and Topaloglu; 2005). This problem is particularly acute in the handling of
22 the state transition, which is typically defined in terms of a transition probability matrix P that specifies
23 the probability that some specific value of the state variable will occur in the next period given the current
24 value of the state and actions variables. The total number of elements of this matrix grows with the square
25 of the number of state values.

26 This note discusses how the curse of dimensionality can be made somewhat less problematic by
27 careful attention to how the transition is handled. In particular it points out that the transition matrix P
28 need not be explicitly defined but instead can be replaced by a function which computes the expectation
29 of future values conditioned on current states and actions. Such a function will be referred to as an

30 expected value (EV) function and its use can have significant advantages both in reducing memory
31 requirements and in speeding up computations (using both function and policy iteration). Two common
32 examples of when such advantages are possible arise when the state variables are conditionally
33 independent or when the transition can be broken into separate phases. The methods discussed in this note
34 are easily implemented using the freely available MATLAB based MDPSolve package (Fackler, 2011)
35 and code for the examples discussed here is available as a supplement.

36 One application area where the curse of dimensionality is particularly problematic is in
37 conservation management of spatial units. For example in Stochastic Patch Occupancy Models (SPOMs)
38 the state variables are binary variables representing the absence or presence of some species on a site
39 (patch). With N sites the state space is 2^N and thus grows exponentially in the number of sites. Another
40 example of such a problem is the reserve site selection problem in which a set of sites are targeted for
41 acquisition by a conservation organization but may instead be acquired and developed for non-
42 conservation uses. In this case each site has three alternative states (available, reserved, developed) and
43 hence the state space is 3^N .

44 This paper first briefly reviews the dynamic programming framework, including a discussion of
45 how index vectors can be used to improve efficiency. It then introduces the concept of an Expected Value
46 (EV) function. Two situations which lead to significant advantages by using EV functions are then
47 discussed and illustrated. The first is the situation in which the state transition occurs in stages, with each
48 stage represented by a sparse transition probability matrix. The second is when a model can be
49 represented in a factored form by a set of conditionally independent state transitions.

50 **Dynamic Programming**

51 The basic components of a DDP model are (1) a reward function $R(S, A)$ which describes the
52 current net benefits of being in a given state S and taking a specified action A , (2) a transition probability
53 matrix, $P(S^+|S, A)$, which gives the transition probability of moving to a specified state S^+ in the next

54 period, given the current state and action and (3) a discount factor $\delta \in [0,1]$ that measures the value of
 55 obtaining a given reward in the next period relative to obtaining it this period. The solution to a dynamic
 56 programming problem is a strategy that defines how the action should be chosen for each value of the
 57 state, $A^*(S)$, and a value function $V(S)$ which describes the value in each state of the sum of the expected
 58 discounted rewards when using the optimal strategy.

59 Standard algorithms for solving dynamic programming problems are based on the Bellman
 60 Equation

$$65 \quad V(S) = \max_A R(S, A) + \delta \sum_{S^+} P(S^+|S, A)V(S^+)$$

61 If there are n_s values of the state variable(s) and n_x possible combinations of state and action values then
 62 V is an n_s element vector, R is an n_x element vector and P is an $n_s \times n_x$ column-stochastic matrix (a
 63 matrix composed of non-negative numbers with columns that sum to 1).² In this case the Bellman
 64 function can be written as

$$73 \quad V = \max_A R_A + \delta P_A^T V^+$$

66 where the A refers to a given strategy. The two standard methods for solving DP problems (function and
 67 policy iteration) both use an initial guess of the vector V and compute the vector $\tilde{V} = R + \delta P^T V$. Each
 68 row of this vector is associated with a specified value for the state and the maximal value for each state
 69 can then be identified. This results in an n_s vector of indices I^a that selects these values of \tilde{V}

$$74 \quad I_i^a = \operatorname{argmax}_{j: I_x(j)=i} \tilde{V}_j$$

70 (the use of index vectors is discussed in more detail in Supplemental Appendix 1). The two methods
 71 differ in how they update V . Function iteration replaces V with $\tilde{V}[I^a]$ whereas policy iteration replaces V
 72 with the solution to the linear system

$$75 \quad (I - \delta P[:, I^a]^T)V = R[I^a].$$

² Alternatively it could be an $n_x \times n_s$ row-stochastic matrix with rows that sum to 1.

76 Both methods repeat this process iteratively until a convergence criterion is met. In general, policy
77 iteration uses fewer iterations but each iteration is more expensive because of the need to perform a linear
78 solve.

79 The state S is typically composed of a set of d_s variables and the size of the state space is the
80 number of possible combinations (tuples) of these variables. A significant challenge in formulating and
81 solving realistic decision models is the so-called curse of dimensionality. The problem size grows
82 exponentially as d_s increases; for example, if all state variables can take on m different values then the
83 size of the state space is m^{d_s} . Of particular importance is that the P matrix can become prohibitively
84 large. Even when sparse (i.e., having many 0 elements) it can use up large amounts of memory and
85 performing the linear solve in policy iteration may become extremely time consuming or even impossible
86 due to memory limitations. Even the matrix-vector operations used to compute $P^T V$ may be prohibitively
87 time-consuming.

88 One approach to rescuing policy iteration which works well for large problems uses iterative
89 linear solvers, including Krylov methods (Barrett et al., 1994). This approach is discussed in Rust (1996)
90 and was demonstrated by Mrkaic (2002) to result in significant reductions in the time required for each
91 iteration when using policy iteration. The use of Krylov methods, such as Generalized Minimum Residual
92 (GMRES) and Bi-Conjugate Gradient-Stabilized (BiCGSTAB), are easily implemented into dynamic
93 programming algorithms in MATLAB because these linear equation solvers are part of the basic
94 MATLAB package.

95 What does not appear to be widely recognized in the literature is the potential for memory and
96 speed efficiencies from not forming the P matrix in the first place. All that is required of function
97 iteration or policy iteration, if a Krylov solver is used, is that $P^T V$ can be evaluated.

98 **Expected Value Functions**

99 An expected value (EV) function produces the same result as $P^T V$ but without the need to
100 explicitly compute P . Specifically, an EV function v transforms the future state vector into its expectation
101 conditional on current states and actions (X):

102
$$v(V^+) = E[V^+ | X]$$

103 An EV function might also use a second input argument,

104
$$v(V^+, I^a) = E[V^+ | X[I^a, :]]$$

105 in which case it is an indexed evaluation that transforms the future state vector into its expectation
106 condition on the states and actions indexed by I^a .

107 The maximization step in the dynamic programming algorithm uses a full EV evaluation:

108
$$I_i^a = \operatorname{argmax}_{j: I_x(j)=i} R_j + \delta[v(V)]_j$$

109 whereas the value function updates use an indexed evaluation. If function iteration is used

110
$$V \leftarrow R[I^a] + \delta v(V, I^a)$$

111 If policy iteration is used then V solves the linear equation:

112
$$h(V) = V - \delta v(V, I^a) = R[I^a]$$

113 Note that this linear solve cannot be performed using direct methods (e.g., LU decomposition) because
114 the matrix operator is not available but can be solved efficiently using iterative Krylov methods.³ Thus
115 both standard methods for solving DP problems are still available when EV functions are used.

116 There are at least two situations in which the use of EV functions is advantageous. The first
117 situation in which large gains are possible with an EV function approach arises when the state transition
118 occurs in phases, $P = P_2 P_1$, where the transition matrix for each phase, P_i , is sparse. Typically P will be
119 far less sparse than its components, in which this case it is possible that $P_1^T (P_2^T V)$ can be computed far

³ The implicit matrix involved here, $I - \delta P^T$, is easily shown to be row-wise strictly diagonally dominant, which is a typical sufficient condition for ensuring that an iterative linear solver converges.

120 faster than $P^T V$ and use far less memory. This will be illustrated with a Stochastic Patch Occupancy
121 Model (SPOM) and with a model in which, in the first stage, the action transforms the state
122 deterministically.

123 The second situation is when two or more sets of the state variables have transition probabilities
124 that are conditionally independent, where conditioning is on subsets of the current state and action
125 variables. Such a situation arises in many dynamic programming models. This is illustrated with a harvest
126 management example and with an SPOM model defined on a network of interconnected sites. To
127 facilitate the specification of such EV functions a set of procedures was developed that allows a user to
128 pass a set of transition matrices for individual state variables, along with information on the conditioning
129 variables involved.

130 **Staged Transitions**

131 The first situation in which there are gains from using the EV function approach arises when the
132 transition can be broken into separate phases, each of which can be described by a sparse transition
133 matrix. Such a situation arises with so-called Stochastic Patch Occupancy Models (SPOMs). Early
134 contributors to this literature are Caswell & Etter (1993), Hanski (1994) and Day & Possingham (1995).
135 In these models there are N sites or patches that can each be classified as either empty or occupied. In one
136 of the phases, the extinction phase, an occupied patch might change to empty with probability e (and if
137 empty it remains so). In the other phase, the colonization phase, an empty patch might change to occupied
138 with probability c (and if occupied it remains so). Typically e and c may differ from patch to patch and
139 will be functions of the current condition of the other patches and of actions that resource managers take.

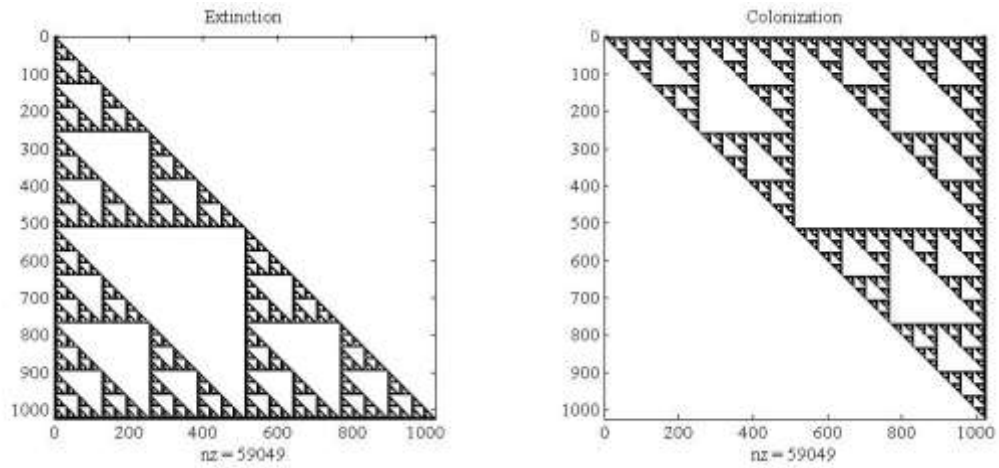
140 In SPOMs the state variable is a vector of N 0s and 1s representing the occupancy status of each
141 patch. The number of possible configurations is 2^N which clearly is a manifestation of the curse of
142 dimensionality. The larger issue for such models, however, is that P has 4^N elements (for any given
143 treatment) and is typically dense or nearly so. The transition matrix however can be decomposed into its

144 extinction and colonization phases, either as $P = EC$ or $P = CE$ where E and C represent the transition
145 probability matrices for the two phases (which order is used depends on when action is taken). For an
146 individual site the site transition matrices for each stage are triangular:

$$147 \quad E_i = \begin{bmatrix} 1 & e_i \\ 0 & 1 - e_i \end{bmatrix} \quad C_i = \begin{bmatrix} 1 - c_i & 0 \\ c_i & 1 \end{bmatrix}$$

148 Note that, in this simple model, the colonization probabilities do not depend on the occupancy status of
149 other patches. The full extinction and colonization transition matrices can therefore be written as a
150 sequence of Kronecker products, e.g., $E = E_1 \otimes E_2 \otimes \dots \otimes E_N$, implying that there are 3^N non-zero
151 values in each of E and C . (this assumes that none of the values of the e_i and c_i are exactly 0 or exactly
152 1), The density of these matrices is thus of $3^N/4^N$ (their sparsity pattern is shown in Figure 1 for $N =$
153 10). Although still problematic, storing 3^N elements in each of two sparse matrices may be feasible for
154 values of N for which storing a dense matrix with 4^N elements is not. Also performing 3^N arithmetic
155 operations twice is much faster than performing 4^N operations once.

156 These results are even more dramatic if each site can be classified into more than 2 categories. If
157 there are m possible categories then there will be m^N values of the state and the transition matrix will
158 contain m^{2N} values. If the two phases represent a decreasing and an increasing phase the single site phase
159 transition matrices will be triangular and thus contain $m(m + 1)/2$ non-zero elements. The number of
160 non-zero elements in full phase transition matrix is this number raised to the power N which implies that
161 the density of the phase transition matrix is $\left(\frac{m+1}{2m}\right)^N$. The density therefore declines towards 2^{-N} as m
162 gets large. Clearly the curse of dimensionality is still present but at least some of its sting has been
163 reduced.



165

166 **Figure 1.** Sparsity patterns for extinction and colonization transition matrices ($N = 10$)

167

168 Table 1 displays the relative times required to do a basic matrix-vector multiplication, which is
169 the basis for Krylov methods, using the full and staged transition approaches. Row 1 displays the time
170 required for 1000 of these operations using the staged form $E^T(C^T V)$ and row 2 shows the same for the
171 full form $P^T V$. At relatively low values of N the full method actually is faster than the staged form, a
172 result that is likely due to the greater efficiency of the matrix multiply operation for full versus sparse
173 formats (this is, of course, dependent on both the software and hardware used). Once N is greater than 10,
174 however, the staged form is faster by an increasingly wide gap, being over 13 times faster for $N = 14$.
175 The third row of the table shows the time required to actually form P by multiplying C and E . This also
176 imposes a significant and avoidable computational burden both in time and memory utilization.

178 **Table 1.** Typical computational times and sparsity for SPOM model

	N						
	8	9	10	11	12	13	14
$E^T(C^T v)$	0.026	0.065	0.086	0.136	0.292	1.672	4.870
Pv	0.014	0.036	0.084	0.801	4.011	15.298	64.277
$P = CE$	0.008	0.008	0.046	0.154	0.724	3.499	19.332
density	0.100	0.075	0.056	0.042	0.032	0.024	0.018

179 Rows 1 & 2 display the time required for 1000 evaluations using the factored form $E^T(C^T v)$ and
 180 full form $P^T v$

181 Row 3 shows the setup time required to a form P

182 Row 4 shows the fraction of non-zero elements in E and C

183

184 Another way that staged transitions can lead to substantial computational gains arises when the
 185 state transition can be written in terms of the so-called post-decision state. For example, in some fisheries
 186 models the future state depends only on escapement which equals the current stock less that harvest. In a
 187 simple model the current stock is the state, the harvest is the action and the escapement is the post-harvest
 188 state.

189 In general if the transition can be divided into a deterministic transition $\tilde{S} = g_1(S, A)$ and a
 190 stochastic transition $S^+ = g_2(\tilde{S}, e)$ then we only require an $n_s \times n_s$ transition matrix P_2 and an n_x index
 191 vector I_1 that defines the g_1 mapping. The expected value function can then be written as $v(V) =$
 192 $[P_2^T V](I_1)$.

193 Conditional Independence

194 Many dynamic models consist of a d_s -element set of state variables that evolve independently
 195 when conditioned on the current state and action variables. The values of the conditioning variables can
 196 organized into an $n_x \times d_x$ matrix X , with each row representing a unique combination of states and
 197 actions. In addition to X a model is defined by a set of d_s conditional probability tables (CPTs), P_i ,
 198 representing the transition probability conditioned on a subset of X and an associated set of index vectors

199 q_i defining the sets of conditioning (parent) variables, with the values of the q_i associated with columns
 200 of X .

201 The simplest case arises when the state variables have disjoint conditioning sets ($q_i \cap q_j = \emptyset$ for
 202 $i \neq j$). In this case the transition matrix can be written as a chain of Kronecker products:

$$203 \quad P = P_1 \otimes \dots \otimes P_{d_s}$$

203 (this was true of the SPOM discussed in the previous section). It is well known that Kronecker product-
 204 vector multiplication can be performed efficiently without actually forming the Kronecker product
 205 (Pereyra and Scherer; 1973). The model of dynamic reserve site selection of Costello and Polasky (2004)
 206 and the harvest management example discussed below both fit this framework.

208 In the more general case, in which the conditioning sets overlap, an EV function can be evaluated
 209 by processing each CPT sequentially using index vectors to define the associated conditioning variables.
 210 The basic approach uses a special indexed multiplication of a 3-D array by a 2-D array:

$$211 \quad y_i(:, k) = y_{i-1}(:, :, I_i^y(k)) P_i(:, I_i^p(k))$$

211 where I_i^y and I_i^p are index vectors that indicate the page (the 3rd dimension) of y_{i-1} and the column of P_i
 212 associated with column k of y_i . Each column of the output y_i is computed as multiplication of an
 213 $(\prod_{j=i+1}^d n_j) \times n_i$ matrix by an n_i vector. At each step the result y_i is reshaped in a 3-D array with n_i
 214 elements in its 2nd dimension. The process is initialized by combining V with P_1 to form y_1 . The
 215 algorithm, which is discussed in greater detail in Supplemental Appendix 2, has the significant advantages
 216 that no copying or shuffling of values in memory is required and that the bulk of the work is performed
 217 using matrix-vector multiplication, which can be implemented in a highly efficient way and uses minimal
 218 memory resources.

220 The number of arithmetic operations is $\sum_{i=1}^d \prod_{j=i}^d n_j \min(k_i, n_s)$ (recall that $n_s = \prod_{j=1}^d n_j$). Contrast this
 221 with an indexed operation using $P[:, I^a]$ which uses n_s^2 arithmetic operations.

222 To illustrate the operations involved consider a problem with 3 state variables and 1 action variable. The
 223 state variable sizes are all n and the action has size n_a . With the action in the last column of X suppose
 224 that the parents vectors are given by

225
$$q_1 = [1 \ 4] \quad q_2 = [1 \ 2 \ 4] \quad q_3 = [2 \ 3 \ 4]$$

226 So future state 1 depends on current state 1 and the action, etc. The EV function is performed in 3 steps
 227 each involving the current intermediate product y_{i-1} and the current CPT P_i . The variables involved
 228 with each array and the number of arithmetic operations required by the indexed multiplication are:

i	y_{i-1}	P_i	# of operations
1	$S_3^+ S_2^+ S_1^+$	$S_1^+ S_1 A$	$n^4 n_a$
2	$S_3^+ S_2^+ S_1 S_2 A$	$S_2^+ S_1 S_2 A$	$n^4 n_a$
3	$S_3^+ S_1 S_2 A$	$S_3^+ S_2 S_3 A$	$n^4 n_a$

229 The total operation count is $3n^4 n_a$. If the full transition matrix is used the operations count is $n^6 n_a$.

230 EV functions can be evaluated using this approach for both full evaluations of the form $v(V)$ and
 231 indexed evaluations of the form $v(V, I^a)$ where I^a is an index vector specifying a strategy. . The latter
 232 form is a bit more complicated to implement and is discussed in detail with an example in Supplemental
 233 Appendix 2.

234 The efficiency of computing an EV function can be influenced both by the sequencing of the
 235 state variables and by performing a preprocessing step in which some of the CPTs are combined to reduce
 236 the amount of computation performed. Determining the optimal sequencing is a difficult problem to solve
 237 and there do not appear to be any polynomial algorithms to solve it. The minimal arithmetic operation
 238 preprocessing of CPTs into groups, however, can be determined using a simple algorithm; this is
 239 discussed in detail in Supplemental Appendix 3.

240 To illustrate the advantage of combining CPTs in a preprocessing step consider 2 CPTs with the
 241 same conditioning sets: $q_1 = [1 \ 2 \ 4]$ and $q_2 = [1 \ 2 \ 4]$. The first two steps with P_1 and P_2 have
 242 operation counts

i	y_i	P_i	# of operations
1	$S_3^+ S_2^+ S_1^+$	$S_1^+ S_1 S_2 A$	$n^5 n_a$
2	$S_3^+ S_2^+ S_1 S_2 A$	$S_2^+ S_1 S_2 A$	$n^4 n_a$

243 If we combine P_1 and P_2 in a preprocessing step to form P_{12} the same operation has

i	y_i	P_{12}	# of operations
1	$S_3^+ S_2^+ S_1^+$	$S_2^+ S_1^+ S_1 S_2 A$	$n^5 n_a$

244 Thus we can do both operations in a single step with the same operation count as the previous first step.

245 **Example: Harvest Management**

246 To demonstrate the extent of the gains consider first the case of managing the harvest of a wild
 247 stock, such as a fishery. Models of this sort go back at least to Clarke & Munro (1975) and many variants
 248 have appeared using both continuous and discrete time formulations. Here we use a fairly simple variant
 249 in which a biological population is commercially harvested with a transition function that can be written
 250 as

$$257 \quad N^+ = f(N, H)u$$

251 where N is the population size, H is the harvest size and u is a random noise term. Suppose that this is
 252 discretized with sorted sets of n_N values of N and n_H values of H . The resulting transition matrix P_N is
 253 $n_N \times n_N n_H$ (this can be viewed as an $1 \times n_H$ vector composed of blocks of size $n_N \times n_N$). In addition the
 254 price received (M) for the harvest evolves dynamically according to

$$258 \quad M^+ = g(M)w$$

255 where w is also a random noise term. Proceeding as before this is discretized and the $n_M \times n_M$ transition
 256 matrix P_M is formed.

259 This is an example in which the conditioning sets (parent variables) form non-overlapping sets
 260 and so the transition matrix can be written as a Kronecker product. If the variables are organized
 261 lexicographically and ordered as (H, N, M) then the combined transition matrix can be written as $P =$
 262 $P_N \otimes P_M$. Rather than using $(P_N \otimes P_M)^T V$ to compute the EV function we can use $P_M^T \check{V} P_N$ where \check{V} is

263 the $n_m \times n_n$ matrix such that $\text{vec}(\tilde{V}) = V$. This expression can be computed as either $P_M^\top(\tilde{V}P_N)$ or
264 $(P_M^\top\tilde{V})P_N$. The former approach requiring $n_m n_n^2 n_a + n_m^2 n_n n_a$ arithmetic operations and the latter
265 requiring $n_m^2 n_n + n_m n_n^2 n_a$; the latter expression therefore unambiguously requires less computational
266 effort.

267 This model was implemented and solved using $n_H = 51$, $n_N = 101$ and $n_M = 101$. The
268 transitions were discretized using linear interpolation weights and either 10000 randomly generated
269 values of u and w (Monte Carlo method) or 21 Gaussian quadrature nodes and weights (quadrature
270 method). The dynamic programming problem was then solved using the full transition matrix with both a
271 direct (LU) linear solver and an iterative Krylov solver (stabilized bi-conjugate gradient) and with 2 EV
272 functions that differed in the order of operations. Using a direct solver required only 6 iterations whereas
273 the use of the Krylov method required 10 iterations (this was true for both discretization methods). The
274 Krylov method typically requires more iterations because it does not attempt to obtain more accuracy
275 than is necessary at each iteration. The optimal decision strategy did not differ between the two linear
276 solve methods.

277 Typical timing results are shown in Table 2. Comparison of the direct and Krylov methods using
278 the full transition matrix (in the first two columns of numbers) clearly demonstrates the advantages
279 possible using Krylov methods rather than direct methods with policy iteration, as has already been
280 demonstrated by Mrkaic (2002). The further advantage of using an EV function is also demonstrated with
281 the better of the two EV functions solving the model approximately 10 times as quickly using Krylov
282 with the full transition matrix and 37-58 times faster than if a direct method is used. The difference in
283 timing results for the two EV functions methods results because the second method performs the first
284 multiplication with P_M which is much smaller than P_N .

285 The differences in the results for the Monte Carlo and the quadrature based methods can be
286 explained by the differences in the degree of sparsity of the transition matrices that the 2 methods

287 produced. P_N and P_M 12% and 39% dense with the Monte Carlo based approach and 20% and 35% with
 288 the quadrature based approach; these values imply densities of 4% and 7% for the full transition matrix.
 289 This leads to a moderate increase in time for the Krylov methods (which rely on simple matrix-vector
 290 operations) and a fairly dramatic increase in time for the direct methods. These results are, of course,
 291 specific to the particular example used here and don't allow the conclusion that the Monte Carlo approach
 292 to discretization should be preferred. Indeed initial computation of the P_N matrices differed dramatically
 293 for the two approaches (3 seconds for the quadrature versus 17 seconds for the Monte Carlo approach).

294

295 **Table 2.** Typical timing results for the harvest management example

discretization approach	solution method			
	full - direct	full - Krylov	$P_M^T(\check{V}P_N)$	$(P_M^T\check{V})P_N$
Monte Carlo	25.76	6.51	2.01	0.69
quadrature	54.73	10.28	2.26	0.95

296

297

298 **Example: Controlling a spatial network**

299 Chadès et al. (2011) developed a Stochastic Patch Occupancy Model (SPOM) for managing
 300 networks of spatial sites that consisted of N sites with an $N \times N$ adjacency matrix C ($C_{ij} = 1$ if sites i and
 301 j are neighbors and 0 otherwise). Each site is either occupied or empty and either treated or not treated:
 302 O/T, O/N, E/T or E/N and a single site can be treated each period.

303 The transition probability for site i depends on whether it is occupied or empty (S_i), treated or not
 304 treated (A_i) and, if empty & not treated, on the # of occupied/untreated neighbors: $q_i = \sum_{j=1}^N C_{ij}S_j (1 -$
 305 $A_j)$. The transition matrix for site i can be represented by a $2 \times (4 + K_i)$ matrix

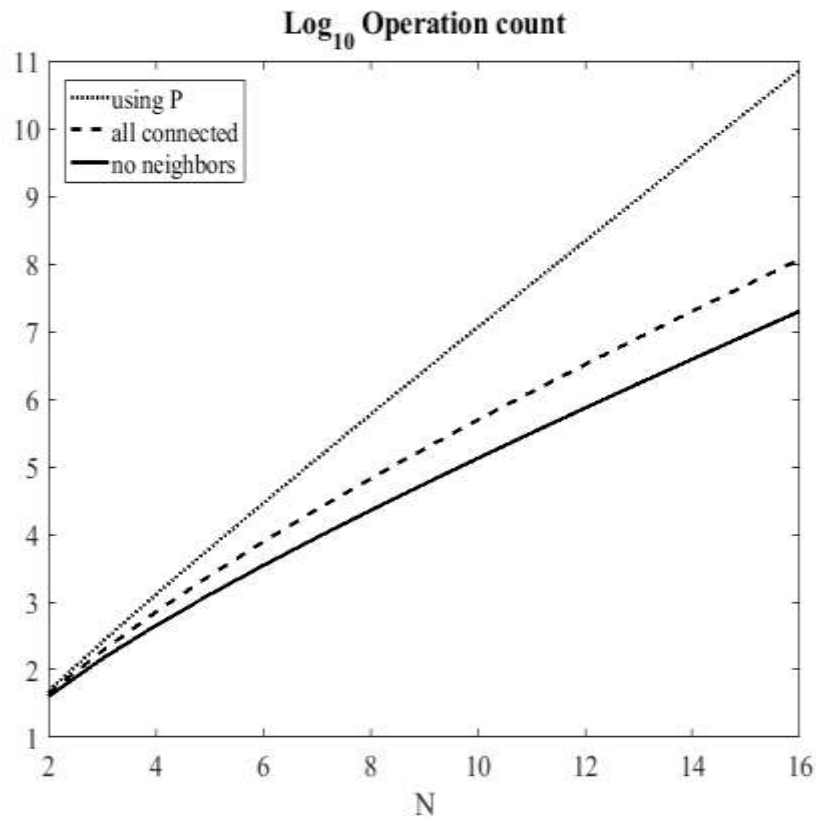
308
$$P_i = \begin{bmatrix} p_{ot} & p_{on} & p_{et} & p_{en}^0 & p_{en}^1 & \dots & p_{en}^{K_i} \\ 1 - p_{ot} & 1 - p_{on} & 1 - p_{et} & 1 - p_{en}^0 & 1 - p_{en}^1 & \dots & 1 - p_{en}^{K_i} \end{bmatrix}$$

306 where the probabilities of occupancy in the next period are p_{ot} (occupied, not treated), p_{on} (occupied,
 307 treated), p_{et} (empty, treated) and p_{en}^j (empty, untreated with j occupied/untreated neighbors, up to K_i).

309 The state space has size 2^N and there are $N + 1$ possible actions (including doing nothing). There are,
310 therefore, $(N + 1)2^N$ state/action combinations

311 If EV functions are used the operation count depends on the density of the network, which can
312 range from all isolated (no neighbors) to all connected, with the operation count increasing as the network
313 becomes more connected. Figure 2 shows the \log_{10} operation count for both isolated and fully connected
314 networks using the EV function approach and compares this to the operation count using the full
315 transition matrix. Even a fully connected network requires significantly fewer operations than using P ;
316 with $N = 16$ there are nearly 3 orders of magnitude fewer operations using the EV function approach.

317 It might seem that, for a fully connected network, there would be no advantage to using an EV
318 function because the transition for each site depends, in principle, on the current state of every other site.
319 In this model, however, the transition for any specific site depends only on how many of its neighbors are
320 occupied. This means that the intermediate factors (the y_i) do not need to grow as fast as they would if
321 the transitions depended on the identities of the occupied neighbors.



323
324 **Figure 2.** Operation count for spatial network model as a function of the number of sites. EV functions
325 are used for the “no neighbors” and “all connected cases.” (SpatNet.m)
326

327

328 Concluding comments

329 This paper introduces the use of expected value (EV) functions as a way to at least partially
330 address curse of dimensionality issues. Although model size still exhibits exponential grow as the number
331 of model variables grows, the use of EV models can nonetheless make feasible the solution of models that
332 might otherwise be out of reach and speed up the solution of models that might previously have been
333 frustratingly slow to solve. This was demonstrated for situations for which the state transition can be
334 broken into separate phases and transitions that can be modeled in factored form.

335 An important challenge for making such an approach more widely used is to recognize when
336 these methods are applicable. Ideally this could be done by the computer so users would not have to

337 engage in complicated programming. In some cases, such as transitions that can be broken into stages, the
338 use of EV functions is fairly natural. It may also be easy to determine if a model can be described in
339 factored form with the state transitions conditioned on subsets of current states and actions. In this case
340 easy-to-use software for creating the EV function has been incorporated into the MDPSolve package.
341 This consists of a function that accepts as inputs the CPTs (P_i), the set of parent variables for each future
342 state variable (q_i) and the matrix of conditioning variables (X) and returns an EV function which can then
343 be passed to the dynamic programming solver.

344 The examples provided here do not cover all of the possible cases for which EV functions may be
345 useful. An important omission is one in which the CPTs for the future state variables are conditioned on
346 noise terms that are common to 2 or more states. Such a noise term cannot be eliminated until all the state
347 variables that it affects are already processed. This typically results in larger intermediate factors, thereby
348 increasing both processing time and memory usage. Nonetheless, a factored approach may still improve
349 on the use of the full transition matrix, especially if there are subsets of state variables which involve
350 nearly disjoint sets of conditioning variables.

351 **References**

- 352 Barrett, R., M. Berry, T.F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine
353 and H. Van der Vorst. 1994. "Templates for the Solution of Linear Systems: Building Blocks for
354 Iterative Methods." SIAM, Philadelphia, PA.
- 355 Caswell, H. & R.J. Etter. 1993. "Ecological interactions in patchy environments, from patch occupancy
356 models to cellular automata." *Lect. Notes Biomath.*, 96: 93–109.
- 357 Chadès, Iadine, Tara G. Martin, Samuel Nicol, Mark A. Burgman, Hugh P. Possingham, and Yvonne M.
358 Buckley. 2011. "General rules for managing and surveying networks of pests, diseases, and
359 endangered species." *PNAS*, 108 (20) 8323-8328. <https://doi.org/10.1073/pnas.1016846108>
- 360 Clark, C. W. & G. R. Munro. 1975. "The economics of fishing and modern capital theory: A simplified
361 approach." *Journal of Environmental Economics and Management*, 2: 92–106.
- 362 Costello, Christopher and Stephen Polasky. 2004. "Dynamic Reserve Site Selection." *Resource and
363 Energy Economics*, 26(2): 157-174.
- 364 Day, J. & H.P. Possingham. 1995. A stochastic metapopulation model with variability in patch size and
365 position. *Theor. Popul. Biol.*, 48: 333–360.
- 366 Fackler, Paul L. 2011. "MDPSolve User's Guide." Available at: <https://sites.google.com/site/mdpsolve/>
- 367 Hanski, I., 1994. "A practical model of metapopulation dynamics." *J. Anim. Ecol.*, 63: 151–162.
- 368 Marescot, Lucile, Guillaume Chapron, Iadine Chadès, Paul L. Fackler, Christophe Duchamp, Eric
369 Marboutin & Olivier Gimenez. 2013. Complex decisions made simple: A primer on stochastic
370 dynamic programming. *Methods in Ecology and Evolution*, 4: 872–884
- 371 Mrkaic, Mico. 2002. "Policy Iteration Accelerated with Krylov Methods." *Journal of Economic
372 Dynamics and Control* 26: 517-45.
- 373 Pereyra, V. & Scherer, G. 1973. Efficient computer manipulation of tensor products with applications to
374 multidimensional approximation. *ACM Transactions Math. Comput.*, 27: 595-605.
- 375 Powell, Warren B. and Huseyin Topaloglu, 2005. Approximate Dynamic Programming for Large-Scale
376 Resource Allocation Problems. *Tutorials in Operations Research*, Chapter X, INFORMS—New
377 Orleans 2005. doi 10.1287/educ.1053.0000
- 378 Puterman, M.L. 1994 *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John
379 Wiley & Sons, New York.
- 380 Rust, John. 1996. "Numerical Dynamic Programming in Economics" in H. Amman, D. Kendrick and J.
381 Rust (eds.) *Handbook of Computational Economics*. Elsevier, North Holland.
- 382 Rust, John. 2008. "Dynamic programming" in Steven N. Durlauf and Lawrence E. Blume (eds.), *The
383 New Palgrave Dictionary of Economics*. Second Edition. Palgrave Macmillan.

384 **Supplemental Appendix 1:**
 385 **Index Vectors**

386 Index vectors are vectors composed of positive integers and can be used for extraction, expansion and
 387 shuffling operations. They are used extensively in matrix based programming environments such
 388 as MATLAB and R. To illustrate let:

$$389 \quad A = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 2 & 0 \\ 2 & 1 \\ 3 & 0 \\ 3 & 1 \end{bmatrix} \quad B = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \\ 2 & 0 & 0 \\ 2 & 0 & 1 \\ 2 & 1 & 0 \\ 2 & 1 & 1 \\ 3 & 0 & 0 \\ 3 & 0 & 1 \\ 3 & 1 & 0 \\ 3 & 1 & 1 \end{bmatrix}$$

390 The index vector $I = [5 \ 6 \ 7 \ 8]$ extracts the rows of B with the first column equal to 2 so $B(I_j, 1) =$
 391 2 for every j . The index vector $I = [1 \ 1 \ 2 \ 2 \ 3 \ 3 \ 4 \ 4 \ 5 \ 5 \ 6 \ 6]$ expands A so $A(I, :) =$
 392 $B(:, [1 \ 2])$. Similarly $I = [1 \ 2 \ 1 \ 2 \ 3 \ 4 \ 3 \ 4 \ 5 \ 6 \ 5 \ 6]$ expands A so $A(I, :) =$
 393 $B(:, [1 \ 3])$. Finally the index vector $I = [1 \ 3 \ 5 \ 2 \ 4 \ 6]$ shuffles the rows of A so they are sorted
 394 by the second column rather than the first:

$$396 \quad A(I, :) = \begin{bmatrix} 1 & 0 \\ 2 & 0 \\ 3 & 0 \\ 1 & 1 \\ 2 & 1 \\ 3 & 1 \end{bmatrix}$$

395
 397 Dynamic programming algorithms can be described in terms of index vectors. Consider a DP model with
 398 2 state variables, each binary, and 3 possible actions
 399

400 The matrix S lists all possible states and X lists all possible state/action combinations:
 401

$$402 \quad S = \begin{bmatrix} 0 & 0 \\ 0 & 1 \\ 1 & 0 \\ 1 & 1 \end{bmatrix} \quad X = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \\ 2 & 0 & 0 \\ 2 & 0 & 1 \\ 2 & 1 & 0 \\ 2 & 1 & 1 \\ 3 & 0 & 0 \\ 3 & 0 & 1 \\ 3 & 1 & 0 \\ 3 & 1 & 1 \end{bmatrix}$$

403
 404 (note that column 1 of X is the action and columns 2 and 3 are the 2 states). The expansion index vector
 405 that gives the states in each row of X is
 406 $I_x = [1 \ 2 \ 3 \ 4 \ 1 \ 2 \ 3 \ 4 \ 1 \ 2 \ 3 \ 4]$

407 This expands S so $S(I_x, :) = X(:, [2 \ 3])$.

408

409 A state dependent strategy can be specified as an extraction index vector with the i th element associated
410 with state i :

411

412 $I^a = [1 \ 6 \ 7 \ 12]$ yields:

414

$$X(I^a, :) = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 0 & 1 \\ 2 & 1 & 0 \\ 3 & 1 & 1 \end{bmatrix}$$

413

415 i.e., a strategy that associates action 1 with state 1, action 2 with states 2 and 3 and action 3 with state 4

416

417 Strategy vectors select a single row of X for each state so $X(I^a, J^s) = S$ where J^s is an index of the
418 columns of X associated with the state variables.

419

420 **Supplemental Appendix 2:**
421 **Computational approach to evaluating EV functions**

422
423 A factored model is defined by a set of d_s conditional transition probability matrices P_i of size $n_i \times m_i$.
424 The computations necessary to compute an EV function can be implemented in a set of d_s multiplication
425 operations involving the CPTs. The multiplication operations have a special form which can be called
426 indexed multiplications. These involve a 3-D array X multiplied by a 2-D array Y with the arrays matched
427 according to 2 index vectors, I^x and I^y , both of length K .

428 The indexed multiplication can be described as follows. Let the inputs X be $m \times n \times p$ and Y be $n \times q$
429 and the output Z be $m \times K$, where $Z_{:k} = X_{::I_k^x} Y_{:I_k^y}$ (the $:$ indicates all elements for a given dimension).

430 Thus each column of the output Z is computed as an ordinary matrix-vector product of one of the pages
431 (3^{rd} dimension) of X and one of the columns of Y . Note that when arrays are stored in column-major form
432 (as is true with MATLAB) the subarrays used in the matrix-vector products are stored in contiguous
433 memory. These matrix-vector products can be computed efficiently with a call to the BLAS gemv
434 procedure (Netlib, BLAS (Basic Linear Algebra Subprograms), <https://www.netlib.org/blas/>). Let this
435 function be represented as $Z = IM(X, I_x, Y, I_y)$. To avoid unnecessary indexing, if the index vector for
436 either X or Y is null (empty) then the index is assumed to equal 1 through K .

437 The algorithm for computing an EV function can now be described. First, set $y_0 = V$ and let y_i be the
438 intermediate product after incorporating the first i CPTs. Let I_i^p and I_i^y be index vectors with length $k_i =$
439 $\prod_{j \in Q_i} n_j$ where $Q_i = \cup_{k=1}^i q_k$. In words, k_i is the size of the space of conditioning variables for the first
440 i state variables.

441 Using the I index vectors a full EV function evaluation is computed using the following pseudo-code:

```
set y = v
reshape y to be  $\prod_{j=2}^d n_j \times n_1$ 
set y  $\leftarrow y * p_1$ 
loop from i = 2 to i = d
    reshape y to be  $(\prod_{j=i+1}^d n_j) \times n_i \times k_{i-1}$ 
    set y  $\leftarrow IM(y, I_i^y, P_i, I_i^p)$ 
return y
```

442

443 The total operation count is $\sum_{i=1}^d p_i k_i$ where $p_i = \prod_{j=i}^d n_j$ is the size of the space of the remaining
444 unprocessed state variables. This can be contrasted to the use of the full transition matrix, which uses
445 $n_s n_x$ operations. Note that variable order matters and ideally we want the k_i to grow slowly. It should
446 also be noted that the reshape operation that transforms a $(\prod_{j=i}^d n_j) \times k_{i-1}$ matrix into a $(\prod_{j=i+1}^d n_j) \times$
447 $n_i \times k_{i-1}$ 3-D array has no computational cost as it does not require access to the elements of the array
448 but merely alters how those elements are interpreted.

449

450 The discussion thus far has applied to a full EV evaluation which returns $E[V(S^+)|X]$ for all state/action
451 combinations. When the dynamic programming algorithm is carried out using policy iteration and Krylov
452 methods most EV evaluations are indexed. Hence we also require an efficient way to compute
453 $E[V(S^+)|X]$ for a specific strategy. A strategy can be defined by the index vector I^a (with length n_s).
454 Although it is possible to simply do a full (non-indexed) evaluation and then extract the elements using I^a
455 such an approach would perform a large amount of unnecessary computations.

456 An alternative uses a set of J_i^p index vectors that expand the columns of P_i to match those of the full X
 457 matrix. Each J_i^p is a vector of length n_x (i.e., equals the # of rows of X). The algorithm could be
 458 initialized as before ($y \leftarrow y * p_1$) and then y could be expanded by setting ($y \leftarrow y(:, J_1^p)$). Then, looping
 459 over the remaining CPTS we could use $y \leftarrow IM(y, [], P_i, J_i^p)$
 460 Where $[]$ represents a null (empty) input. A more efficient approach recognizes that early in the operation
 461 it is generally more efficient to use the I^p indices and latter it is more efficient to use the J^p indices. At
 462 some point the length of I^p is greater than n_s (the length of I^a), at which point it would be more efficient
 463 to switch to the use of the J^p indices. To implement this we also need an additional index vector J^y to
 464 expand y_i at the time the switch is made.

465 The indexed EV function evaluation is described by the following pseudo-code:

```

set y = v
reshape y to be  $\prod_{j=2}^d n_j \times n_1$ 
set y  $\leftarrow y * p_1$ 
set useI = true
loop from i = 2 to i = d
  if  $m_i > n_s$ 
    reshape y to be  $(\prod_{j=i+1}^d n_j) \times n_i \times m_{i-1}$ 
    and expand  $y(:, :, k) \leftarrow y(:, :, J^y(I^a(k)))$ 
    set useI = false
  if useI=true
    reshape y to be  $(\prod_{j=i+1}^d n_j) \times n_i \times m_{i-1}$ 
    set y  $\leftarrow IM(y, I_i^y, P_i, I_i^p)$ 
  otherwise
    set y  $\leftarrow IM(y, [], P_i, J_i^p(I^a))$ 

```

466

467 To illustrate the impact of this algorithm recall the numerical example given in the paper. Furthermore,
 468 suppose that $n < n_a < n^2$ and note that a strategy index has length $n_s = n^3$. The I_i indices have sizes
 469 nn_a, n^2n_a and n^2n_a . The crossover from I to J indexing would therefore occur in step 2.

i	y_i	P_i	# of operations
1	$S_3^+ S_2^+ S_1^+$	$S_1^+ S_1 A$	$n^4 n_a$
2	$S_3^+ S_2^+ S_1 S_2 S_3$	$S_2^+ S_1 S_2 S_3$	n^5
3	$S_3^+ S_1 S_2 S_3$	$S_3^+ S_1 S_2 S_3$	n^4

470 The total operation count is $n^4(n_a + n + 1)$. If the full transition matrix is used by extracting the
 471 appropriate columns of P : $P[:, I^a]$ the operation requires n^6 operations.

472

473

474 **Supplemental Appendix 3:**
475 **Optimal preprocessing of CPTs**

476 It can be advantageous to preprocess groups of state variables into joint CPTs, especially when the
477 variables in the group have similar sets of conditioning variables. The optimal grouping of operations can
478 be solved using an $O(d^3)$ dynamic programming algorithm that is similar to the approach used to address
479 the well-known matrix chain multiplication problem. Given a variable order the cost of incorporating a
480 CPT that groups variables i through $j \geq i$ is $C_{ij} = p_i m_j$, where $p_i = \prod_{k=i}^d n_k$ and m_j is the number of
481 tuples of the parents of variables 1 through j . For each (i, j) we can evaluate whether breaking the
482 grouped variables into two further groups results in a less costly set of operations:

$$485 \quad M_{ij} = \min \left(C_{ij}, \min_{k \in \{0, \dots, j-i+1\}} M_{i, i+k} + M_{i+k+1, j} \right)$$

483 The minimal cost grouping is given by M_{1d} . This is optimal for a full evaluation. For an indexed
484 evaluation we could instead define

$$486 \quad C_{ij} = p_i \min(m_j, n_s)$$

487 By storing where splits occur the optimal groupings can be determined.