A Online Appendix

To Accompany "Multi-stage Monte Carlo Method for Solving Influence Diagrams Using Local Computation,"

by John M. Charnes and Prakash P. Shenoy

A.1 Background

As our multi-stage Monte Carlo method requires finding the maxima of correlated random variables generated through sampling, we begin by discussing a lemma that appears as an exercise in David [1981, pp. 51–52].

Lemma 1 *If* (X, Y) *has the bivariate normal* $\mathrm{N}(0, 0, 1, 1, \rho)$ *distribution, then*

$$\mathrm{E}[\mathrm{Max}(X, Y)] = \sqrt{\frac{1-\rho}{\pi}} \textit{ and } \mathrm{Var}[\mathrm{Max}(X, Y)] = 1 - \frac{1-\rho}{\pi}.$$

We would like to generalize the statements of Lemma 1 in several ways.

Note that the distribution of $M_0 = \mathrm{Max}(X, Y)$ as defined in Lemma 1 is non-normal, with a probability density function (pdf) that is bell-shaped but skewed slightly to the right.

Now consider the distribution of $M_1 = \mathrm{Max}(X, Y)$, where (X, Y) has the bivariate normal $\mathrm{N}(0, a, 1, 1, \rho)$ distribution, for which we assume that $a > 0$ without loss of generality. As a increases, the pdf of M_1 more closely resembles the pdf of Y. Thus, for large a, the random variable M_1 is approximately normally distributed with mean a and variance 1.

Because the variance is a scale parameter for the normal distribution, if $M_2 = \mathrm{Max}(X, Y)$, where (X, Y) has the bivariate normal $\mathrm{N}(0, 0, \sigma^2, \sigma^2, \rho)$ distribution, then

$$\mathrm{Var}[M_2] = \sigma^2 \left[1 - \frac{1-\rho}{\pi}\right]. \tag{A.1}$$

Finally, although Lemma 1 does not generalize easily to the case of more than two random variables, results provided by Afonja [1972] can be used to obtain an upper bound on the variance of the maximum of more than two elements of a multivariate normal distribution. Specifically, if $M_3 = \mathrm{Max}(X_1, X_2, \ldots, X_n)$, where $(X_1, X_2, \ldots, X_n)$ has the multivariate normal distribution with all means equal to zero, all variances equal to one and all covariances equal to ρ, then

$$\mathrm{Var}\,[M_3] \le 1 - \frac{1-\rho}{\pi}. \tag{A.2}$$

A.2 Canonical Example

Consider a decision problem consisting of two decision variables, A and C, and two chance variables, B and D, depicted by the influence diagram in Figure A.1, and by the decision tree in Figure A.2. We assume all variables have two states. The joint utility function factors additively into two factors u and v. With the Multi-stage Monte Carlo method, we use sampling in Stage 1 to form estimates $Z_{bc}, Z_{bc'}, Z_{b'c}$, and $Z_{b'c'}$ of the mean utilities associated with D and estimates Y_b and $Y_{b'}$ of the mean utilities associated with C (see Figure A.2). In Stage 2, we use sampling to form estimates X_a and $X_{a'}$ of the mean utilities associated with B and determine estimate W of the mean utility associated with A.

We wish to place a bound on $\mathrm{Var}[W]$, the variance of the estimated mean utility at A. At each stage, we can reduce the variance of the estimates to an arbitrarily low value by running the simulation for a sufficiently large number of iterations. Let σ_1^2 and σ_2^2 denote the upper bound on the conditional variance of each estimate obtained in Stages 1 and 2, respectively. In this appendix, we show that $\mathrm{Var}[W] \le \sigma_1^2 + \sigma_2^2$. After all sampling is complete, we use the standard normal cumulative distribution function $\Phi(\cdot)$ to form the confidence interval $W \pm \Phi^{-1}(1-\alpha/2)\sqrt{\mathrm{Var}[W]}$. If the half-width of this interval is smaller than ϵ, we are confident that we have obtained an approximate (ϵ, α)

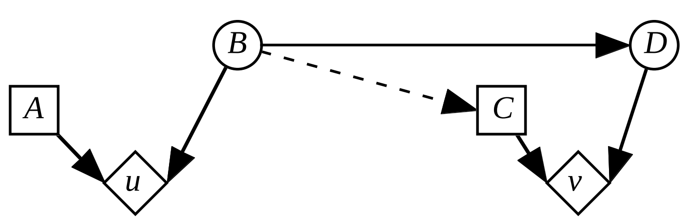

Figure A.1: Influence diagram representation of the canonical example.

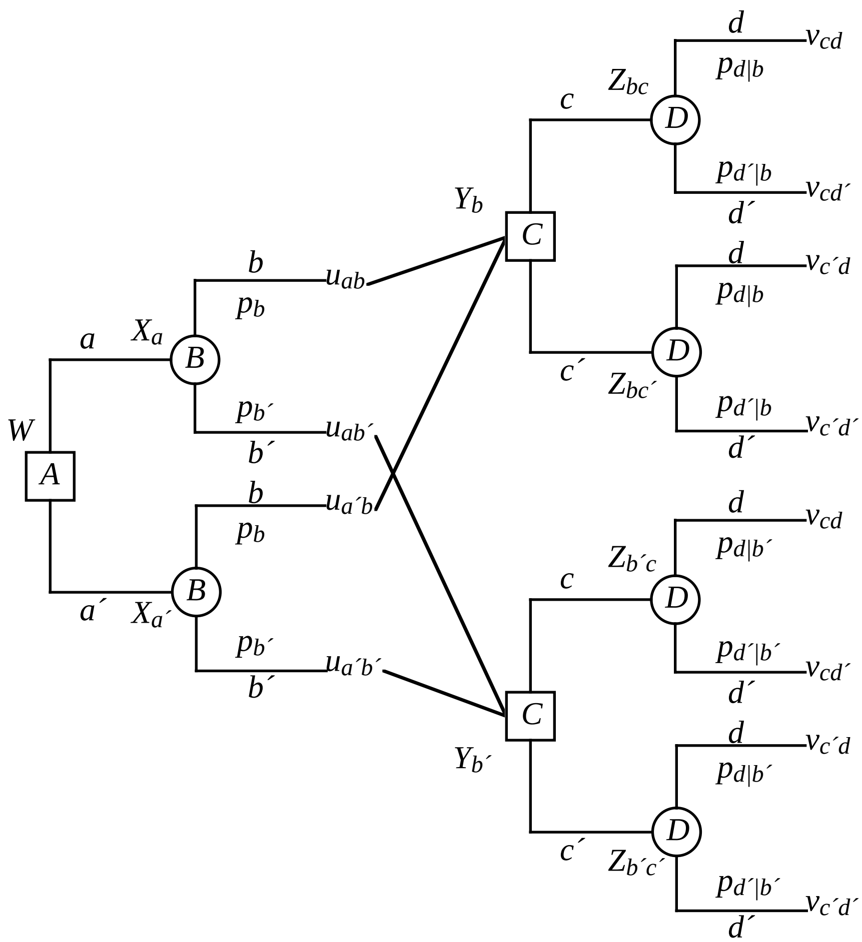

Figure A.2: Decision tree representation of the canonical example.

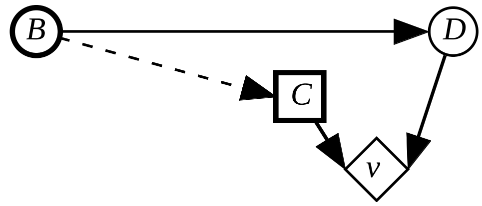

Figure A.3: Stage 1 Influence Diagram.

estimate of mean utility. In practice we find our method to be conservative in the sense that after many replications of the method, more than $100(1-\alpha)\%$ of the estimates are within ϵ of the true value.

A.3 Justification

In Stage 1 of the canonical example (see Figure A.3), for each state of B and C, we sample chance variable D repeatedly (using each of its conditional distributions given B) and estimate the expected utilities from utility function v. We denote the estimated values by random variables $Z_{bc}, Z_{bc'}, Z_{b'c}$, and $Z_{b'c'}$. Notice that these four random variables are mutually independent by virtue of the random numbers used in the simulation. By the Central Limit Theorem, these four random variables are approximately normal with expected values $\mu_{Z_{bc}}, \mu_{Z_{bc'}}, \mu_{Z_{b'c}}$, and $\mu_{Z_{b'c'}}$, respectively, where

$$\begin{aligned}
\mu_{Z_{bc}} &= p_{d|b}v_{cd} + p_{d'|b}v_{cd'}, \\
\mu_{Z_{bc'}} &= p_{d|b}v_{c'd} + p_{d'|b}v_{c'd'}, \\
\mu_{Z_{b'c}} &= p_{d|b'}v_{cd} + p_{d'|b'}v_{cd'}, \text{ and} \\
\mu_{Z_{b'c'}} &= p_{d|b'}v_{c'd} + p_{d'|b'}v_{c'd'}.
\end{aligned}$$

In Stage 1, we do sufficient sampling to ensure that

$$\text{Max}\,\{\text{Var}[Z_{bc}], \text{Var}[Z_{bc'}], \text{Var}[Z_{b'c}], \text{Var}[Z_{b'c'}]\} \le \sigma_1^2.$$

We complete Stage 1 by determining a decision function for C by choosing $Y_b = \text{Max}\{Z_{bc}, Z_{bc'}\}$, and $Y_{b'} = \text{Max}\{Z_{b'c}, Z_{b'c'}\}$. Using independence and Lemma 1, one can show that $\text{Var}[Y_b] \le (1-1/\pi)\sigma_1^2$, and $\text{Var}[Y_{b'}] \le (1-1/\pi)\sigma_1^2$. Thus, $\text{Var}[Y_b] \le \sigma_1^2$ and $\text{Var}[Y_{b'}] \le \sigma_1^2$. Let y_b and $y_{b'}$ denote the realizations of Y_b and $Y_{b'}$ respectively, obtained in Stage 1. The values y_b and $y_{b'}$ are used to form the estimated utility function $\hat{v}$ for use in Stage 2.

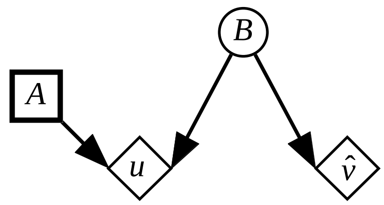

Figure A.4: Stage 2 Influence Diagram.

Next we move to Stage 2 (see Figure A.4), in which for each state of decision variable A we sample for chance variable B (using its distribution), look up the corresponding values of utility factor u and the estimated utility function $\hat{v}$, and compute the estimated values X_a and $X_{a'}$. By the Central Limit Theorem, given $(Y_b, Y_{b'}) = (y_b, y_{b'})$, X_a and $X_{a'}$ are approximately normal with expected values $\mu_{X_a} = (p_b u_{ab} + p_{b'} u_{ab'}) + (p_b y_b + p_{b'} y_{b'})$ and $\mu_{X_{a'}} = (p_b u_{a'b} + p_{b'} u_{a'b'}) + (p_b y_b + p_{b'} y_{b'})$, respectively. Notice that given $(y_b, y_{b'}), X_a$ and $X_{a'}$ are conditionally independent by virtue of the random numbers used in the simulation. At Stage 2, we do sufficient sampling to ensure that

$$\text{Max}\,\{\text{Var}[X_a|y_b, y_{b'}], \text{Var}[X_{a'}|y_b, y_{b'}]\} \le \sigma_2^2.$$

Finally, we determine a decision function for A by estimating $W = \text{Max}\,\{X_a, X_{a'}\}$. In the next subsection, we show that $\text{Var}[W] \le \sigma_1^2 + \sigma_2^2$.

A.4 Bayes Net

Note that the joint distribution of Z_{bc}, $Z_{bc'}$, $Z_{b'c}$, $Z_{b'c'}$, Y_b, $Y_{b'}$, X_a, $X_{a'}$, and W can be represented by a Bayes net as shown in Figure A.5, where Y_b, $Y_{b'}$, and W are conditionally deterministic variables depicted by double-bordered ellipses. The Bayes net provides another view of the stochastic behavior of our method. The estimators X_a and $X_{a'}$ are dependent random variables, but they are conditionally independent given Y_b and $Y_{b'}$.

In Stage 1, we sample sufficiently so that

$$\text{Max}\,\{\text{Var}[Z_{bc}], \text{Var}[Z_{bc'}], \text{Var}[Z_{b'c}], \text{Var}[Z_{b'c'}]\} \le \sigma_1^2.$$

Using the result in expression (A.1),

$$\text{Var}[Y_b] \le (1 - 1/\pi)\sigma_1^2, \text{ and } \text{Var}[Y_{b'}] \le (1 - 1/\pi)\sigma_1^2. \tag{A.3}$$

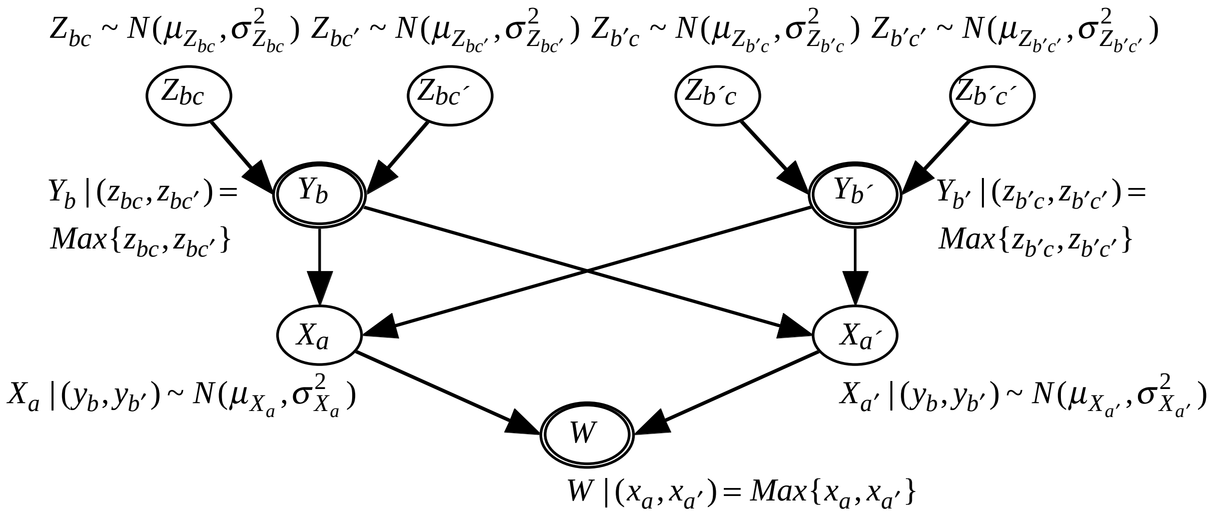

Figure A.5: Bayes net representation of the relationships among the random variables in the canonical decision problem.

In Stage 2, we sample sufficiently so that

$$\mathrm{Max}\,\{\mathrm{Var}[X_a|y_b, y_{b'}], \mathrm{Var}[X_{a'}|y_b, y_{b'}]\} \le \sigma_2^2.$$

From a property of conditional variance,

$$\begin{aligned} \mathrm{Var}[X_a] &= E\left[\mathrm{Var}[X_a|y_b, y_{b'}]\right] + \mathrm{Var}\left[E[X_a|y_b, y_{b'}]\right] \\ &\le \sigma_2^2 + \mathrm{Var}[(p_b u_{ab} + p_{b'} u_{ab'}) + (p_b Y_b + p_{b'} Y_{b'})] \\ &= \sigma_2^2 + p_b^2 \mathrm{Var}[Y_b] + p_{b'}^2 \mathrm{Var}[Y_{b'}] \\ &\le \sigma_2^2 + (1 - 1/\pi)\sigma_1^2. \end{aligned} \tag{A.4}$$

Similarly, $\mathrm{Var}[X_{a'}] \le \sigma_2^2 + (1 - 1/\pi)\sigma_1^2$.

In moving from the second to the third lines in expression (A.4), we were able to assume that Y_b and $Y_{b'}$ are independent from the Bayes net in Figure A.5. However, with more stages, where this independence might not necessarily hold, expression (A.4) is still valid. Denote $\rho = \mathrm{Corr}(Y_b, Y_{b'})$. Then expression (A.4) becomes

$$\begin{aligned} \mathrm{Var}[p_b Y_b + p_{b'} Y_{b'}] &= p_b^2 \mathrm{Var}[Y_b] + p_{b'}^2 \mathrm{Var}[Y_{b'}] + 2 p_b p_{b'} \rho \sqrt{\mathrm{Var}[Y_b] \cdot \mathrm{Var}[Y_{b'}]} \\ &\le \left(p_b \sqrt{\mathrm{Var}[Y_b]} + p_{b'} \sqrt{\mathrm{Var}[Y_{b'}]}\right)^2 \\ &\le (1 - 1/\pi)\sigma_1^2 (p_b + p_{b'})^2 \\ &= (1 - 1/\pi)\sigma_1^2 \end{aligned}$$

from expression (A.3) and the facts that $p_b + p_{b'} = 1$, and $\rho \le 1$.

Notice that although Y_b and $Y_{b'}$ are not normally distributed, their probability density functions are bell-shaped but skewed slightly to the right as was the pdf of the random variable M_0 discussed earlier. (The skewness decreases with increases in $|\mu_{Z_{bc}} - \mu_{Z_{bc'}}|$ and $|\mu_{Z_{b'c}} - \mu_{Z_{b'c'}}|$). Thus, we regard Y_b and $Y_{b'}$ as approximately normal, independent random variables. Also, X_a and $X_{a'}$ are conditionally independent, approximately normal random variables (by the Central Limit Theorem). Thus it follows from the theory of multivariate normal distributions that the marginal distributions of X_a and $X_{a'}$ are normal with some correlation ρ. Appealing to expression(A.1), we get

$$\mathrm{Var}[W] \leq \left[\sigma_2^2 + (1 - 1/\pi)\sigma_1^2\right](1 - (1 - \rho)/\pi) \leq \sigma_1^2 + \sigma_2^2, \tag{A.5}$$

as desired.

In the canonical example, we have assumed two alternatives for each decision variable. With more than two alternatives, although Lemma 1 does not generalize easily, our result (A.5) appears to hold based on numerical integration, simulation experiments, and the bound on the variance of the random variable M_3 reported in expression (A.2).

In the canonical example we have only two stages. For problems having $k > 2$ stages, an analogous bound is

$$\mathrm{Var}[W] \leq \sigma_1^2 + \cdots + \sigma_k^2,$$

where σ_j^2 is the maximum of the conditional variances in stage j, for $j = 1, \ldots, k$. This can be seen easily by induction using (A.5), where σ_1^2 represents $\sigma_1^2 + \cdots + \sigma_{k-1}^2$ and σ_2^2 represents σ_k^2.

References

1. Afonja, B., "The Moments of the Maximum of Correlated Normal and t-Variates," *Journal of the Royal Statistical Society. Series B*, Vol. 24, No. 2 (1972), pp. 251–262.

2. David, H. A., *Order Statistics*, second edition. New York: Wiley, 1981.