Contemporary Mathematics

When to Stop Consulting

D. Marc Kilgour and Steven J. Brams

Abstract. Consider a decision-maker whose objective is to maximize the probability of making a correct binary (Yes-No) decision. To gain information, the decision-maker can consult independent experts. But how many? When should consultation stop?

We model the decision-maker's current state of knowledge using the number of recommendations for each alternative. After every consultation, the decision-maker decides whether to make a decision immediately, based on the advice already received, or to consult again. We impose a ceiling on the number of consultations and, possibly, a cost for each consultation.

Surprisingly, we find that even when consultation is free, it can be optimal *not* to consult. Less surprising, consultation may be uniquely optimal when evidence is scarce, especially if it is approximately balanced between Yes and No. For a decision-maker who, initially, is maximally uncertain, we obtain complete optimal strategies when consultation is free, and we show that there is a fundamental limit on expected utility ($\frac{3}{4}$, where the minimum utility is 0 and the maximum is 1), even when consultation is unlimited. We then explore the strategic effects of a positive cost, linking them to the zero-cost case. We also compare our results on consultation to strategies in the Secretary Problem, another decision problem in which information is collected sequentially until a decision is made.

Individuals facing difficult decisions often ask for advice from family, friends, or colleagues. New points of view and different expertise can provide useful facts and relevant perspectives. But consulting is not without costs, including the cost of delaying the eventual decision. Our concern is whether and when consultation stops helping. At what point does consulting lose its effectiveness as a strategy? How much should a decision-maker consult?

The decision to stop consulting can be crucial, whether the decision-maker is a powerful political leader or a private individual with limited scope for action. For example, as the Cuban Missile Crisis developed, U.S. President John F. Kennedy consulted individually with members of EXCOMM, the Executive Committee of the

1991 *Mathematics Subject Classification.* Primary 90B50; Secondary 62C05.

The first author was supported in part by NSERC Discovery Grant #8588.

National Security Council, before deciding on a response. [**A, B, C**] Anecdotal evidence suggests that the committee 'hawks' dominated the early recommendations, but Kennedy kept on consulting. A prominent early 'dove,' U.N. Ambassador Adlai Stevenson, was later joined by others; eventually, opposition to escalation prevailed.

We model the decision-maker's problem, focusing on strategies for deciding when to stop consulting. We also ask how limitations on the amount of consultation affect the quality of the (eventual) decision, and how much consultants can help. Our objective is to provide advice to decision-makers about when to stop consulting and get on with the decision.

1. Introduction

This paper concerns a decision-maker (DM) who must make a binary (Yes/No) decision. One of these two alternatives is correct, but DM does not know which one. DM's objective is to maximize the probability of making the correct choice. We do not specify the details, but instead assume that, shortly after the choice is made, it becomes clear whether or not it was correct. We model DM as receiving 1 unit of utility for a correct choice, and 0 otherwise; we ask whether there is a consultation strategy that maximizes expected utility.

We assume that DM has beliefs about which choice is correct, which we capture in DM's (subjective) probability, p, that Yes is correct. DM will use any advice in favor of Yes or No to revise appropriately the value of p.

We model consultation as an iterative process in which, at any time, DM can choose to Stop—make a choice of Yes or No based on the current value of p—or to Consult, obtaining advice from a new "expert." After consultation, DM adjusts the value of p to reflect the advice received, and then chooses whether to Stop or to Consult again.

A DM who chooses Stop makes a Yes/No decision immediately. It is clear that, to maximize DM's expected utility, the best choice X_{op} depends on the current value of p; specifically,

$$X_{op} = \begin{cases} \text{Yes,} & \text{if } p \geq \frac{1}{2} \\ \text{No,} & \text{if } p \leq \frac{1}{2} \end{cases}$$

We propose two restrictions on the amount of consultation:

- DM may Consult a maximum of n times
- DM's final utility is reduced by c for each consultation

Thus our model contains two parameters: $n \geq 0$, an integer representing the ceiling on the number of consultations; and $c \geq 0$, representing the cost (including delay cost) of each consultation.

Before we describe the details of our model, we compare it to another sequential decision problem, the Secretary Problem. [**D, E, F**] Both the Secretary Problem and the Consultation Problem involve sequences of information; the problem ends when a decision is made, which may occur at any point in the sequence. In the Secretary Problem, there is a sequence of job candidates, and the decision-maker must select one; the only available information about a candidate is their ordinal ranking relative to previous candidates. The decision to hire a candidate must be made before the next candidate is interviewed. Traditionally, success is measured by the probability that the candidate chosen is, in fact, the overall best candidate. In the Consultation Problem, the alternatives are fixed but estimates of their utilities change as evidence is accumulated (utilities are also diminished by the cost of consultation, if $c > 0$). The Secretary Problem and the Consultation Problem are different, but both concern if and when a finite sequence of random events should trigger a decision.

2. Model

We assume that (1) initially, DM is maximally uncertain, and (2) DM sees each possible consultant as equally credible. (Possible variations in these assumptions are discussed in the Conclusions.) We implement our assumptions by modeling DM's beliefs about p as DM's **state**, (h, k), where h and k are positive integers. In state (h, k), DM's probability that Yes is correct is $p = \frac{h}{h+k}$.

We assume that DM's initial state is $(h, k) = (1, 1)$, so DM's initial probability that Yes is correct is $\frac{1}{2}$. Let $h - 1$ be the number of Yes recommendations that DM has received so far, and let $k - 1$ be the number of No recommendations. We further assume that, in state (h, k), DM believes that the probability that the next consultant will recommend Yes is $\frac{h}{h+k}$, which is identical to DM's current probability that Yes is correct. Thus, a consultation implies a state transition

$$(h, k) \longrightarrow \begin{cases} (h+1, k) & \text{with probability } \frac{h}{h+k} \\ (h, k+1) & \text{with probability } \frac{k}{h+k} \end{cases}$$

For example, suppose that DM consults 3 times and receives 1 Yes and 2 No's. Then after these consultations, DM's state is $(h, k) = (2, 3)$, DM's subjective probability that Yes is correct is $\frac{2}{5}$, which is also DM's subjective probability that the next consultant will recommend Yes, in which case the state would transition to $(3, 3)$.

Given a decision to Stop in state (h, k), DM maximizes utility by choosing

$$X_{op} = \begin{cases} \text{Yes} & \text{if } h \geq k \\ \text{No} & \text{if } h \leq k \end{cases}$$

Thus, at state (h, k), DM's immediate expected utility for choosing Stop is

$$u(X_{op}) = u(h, k) = \frac{\max\{h, k\}}{h + k}. \tag{2.1}$$

But in state (h, k), there have been $h + k - 2$ consultations, so DM's net utility will equal $u(X_{op}) - (h + k - 2)c$. For instance, in state $(h, k) = (2, 3)$, DM can obtain expected utility of $\frac{3}{5} - 3c$ by choosing No.

DM's current state (h, k) reflects the recommendations has received. The set of possible states forms a simple lattice, shown in Figure 1, which also shows the transition probabilities moving to the right, starting at state $(h, k) = (1, 1)$. Figure 1 also shows DM's immediate expected utility for choosing Stop at each state (h, k), as determined by (2.1).

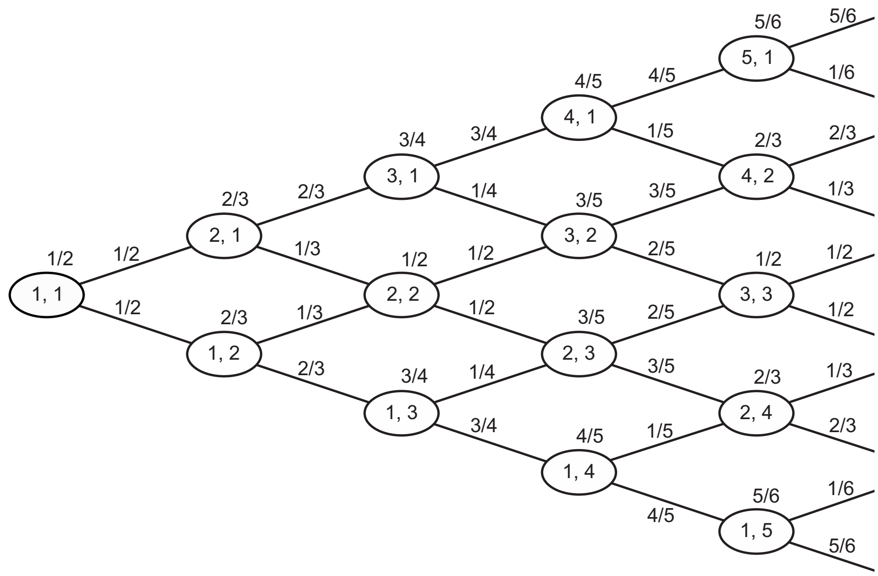

FIGURE 1. States, transition probabilities, and immediate utilities

In fact, assuming consultation cost $c = 0$, our model is a standard Bayesian decision model in which, progressing through the decision tree, DM's current probabilities and expected payoffs reflect DM's experience so far. Thus, DM's probability for p is Beta-distributed [**G**], with probability density function

$$f(x) = \frac{\Gamma(h+k)}{\Gamma(h)\Gamma(k)} x^{h-1}(1-x)^{k-1} \quad (0 < x < 1)$$

(If ℓ is a positive integer, then $\Gamma(\ell) = (\ell - 1)!$.) In state (h, k), the expected value of p is

$$E[p] = \frac{h}{h+k}$$

so that DM's current belief about p is equal to the mean of the current distribution of p.

3. No-Cost Models

We begin by solving all No-Cost Models, in which $c = 0$ and $n = 0, 1, \ldots$. If the process arrives at a state (shown below as a node), DM must decide whether to Stop (S) or Consult (C)—except, of course, that S is the only option if Consult has already been chosen n times. DM's expected utility for choosing optimally at node $(h, k) = (1, 1)$ and, if appropriate, all subsequent nodes, is denoted $U_c(n)$. We first study $U_0(n)$.

3.1. No-Cost Model: $n = 0$**.** If no consultations are allowed, DM must Stop at $(h, k) = (1, 1)$. There are no decisions to make, and by (2.1) DM's expected utility is $U_0(0) = \frac{1}{2}$.

3.2. No-Cost Model: $n = 1$**.** If exactly one consultation is allowed, DM may stop at $(h, k) = (1, 1)$, which would produce expected utility $U_0(0) = \frac{1}{2}$, or consult once, ending up at either $(2, 1)$ or $(1, 2)$. Using (2.1), DM's decision problem is as shown in Figure 2, where nodes labelled '0' are random (i.e., choices by Nature) and thicker (darker) branches indicate choices consistent with backward induction from the endpoints.

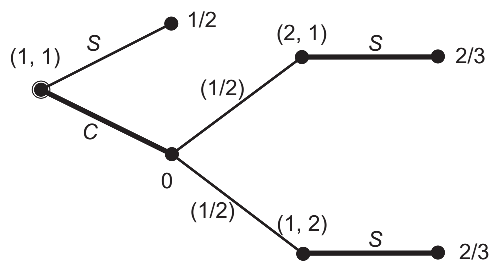

FIGURE 2. Decision Problem, $c = 0$, $n = 1$

It is easy to see that DM's optimal choice at $(1, 1)$ is C, and DM's expected utility for choosing optimally throughout is $U_0(1) = \frac{2}{3}$.

3.3. No-Cost Model: $n = 2$**.** If exactly two consultations are allowed, DM has a decision to make at $(h, k) = (1, 1)$; if that decision is Consult, then there is another decision at $(2, 1)$ or $(1, 2)$. Using (2.1), DM's decision problem is as shown in Figure 3. Note that the figure is not a conventional decision tree, as the node at $(2, 2)$ has indegree 2. But it is easy to see that this simplification does not change the analysis.

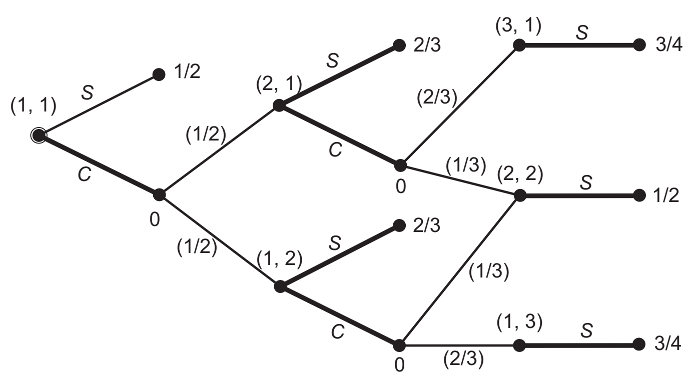

FIGURE 3. Decision Problem, $c = 0$, $n = 2$

Again, the darker branches in Figure 3 indicate rational choices. It is clear that at $(1, 1)$ DM's optimal choice is C, and at $(2, 1)$ and $(1, 2)$ it is C or S. DM's expected utility starting at $(1, 1)$ is $U_0(2) = \frac{2}{3}$.

3.4. No-Cost Model: $n = 3$. By a similar analysis, it can be shown that, if $c = 0$ and $n = 3$, DM's optimal choice at $(1,1), (2,1), (1,2)$ and $(2,2)$ is C, whereas at $(3,1)$ and $(1,3)$ it is either C or S. The expected utility of the optimal strategy $U_0(3) = \frac{7}{10}$.

3.5. No-Cost Model: General Analysis. Assuming that the current state is (h,k), we compare two options, (S) Stop Immediately versus (CS) Consult and then Stop Immediately. First note that, at state (h,k), there have already been $h+k-2$ consultations, so CS is infeasible unless $n > h+k-2$.

Lemma 3.1. *Assume $c = 0$ and $n > h+k-2$. If $h \neq k$, then at (h,k) DM's expected utilities for option S and option CS are equal.*

Proof. Assume $h > k$. At (h,k), S yields expected utility $\frac{h}{h+k}$. Option C produces state $(h+1,k)$ with probability $\frac{h}{h+k}$ and state $(h,k+1)$ with probability $\frac{k}{h+k}$. Using (2.1), CS yields expected utility

$$\frac{h}{h+k} \cdot \frac{h+1}{h+k+1} + \frac{k}{h+k} \cdot \frac{h}{h+k+1} = \frac{h^2+h+hk}{(h+k)(h+k+1)} = \frac{h}{h+k}.$$

The proof is analogous if $h < k$. □

The message of Lemma 3.1 is that it is never an advantage to Consult if you plan to Stop immediately after consulting. Crucial to this result is that the state, (h,k) is not located on the centre-line of Figure 1—in other words, that $h \neq k$.

Along the centre-line of Figure 1, the situation is different, as shown by the next lemma.

Lemma 3.2. *Assume $c = 0$ and $n > 2h-2$. At state (h,h), option CS yields greater expected utility than option S.*

Proof. At (h,h), S yields expected utility $\frac{1}{2}$ by (2.1), while option C produces state $(h+1,h)$ with probability $\frac{1}{2}$ and state $(h,h+1)$ with probability $\frac{1}{2}$. Again using (2.1), CS at (h,h) yields expected utility

$$\frac{1}{2} \cdot \frac{h+1}{2h+1} + \frac{1}{2} \cdot \frac{h+1}{2h+1} = \frac{h+1}{2h+1} > \frac{1}{2}.$$

□

Lemma 3.2 provides an important contrast with Lemma 3.1. If it happens that $h = k$, then choosing CS, Consult and then Stop, is an improvement on choosing S, Stop immediately.

Lemma 3.3. *Assume $c = 0$ and $n > h+k-2$. If either S does not maximize DM's expected utility at $(h+1,k)$ or S does not DM's maximize expected utility at $(h,k+1)$, then only C maximizes DM's expected utility at (h,k).*

Proof. If $h = k$, Lemma 3.2 shows that Stop at (h,k) cannot maximize expected utility. If $h \neq k$, Lemma 3.1 shows that S and CS produce equal expected utility. Suppose that, at $(h+1,k)$, DM has a choice that begins with C and maximizes expected utility. Then at (h,k) the choice of C must maximize expected utility. The argument is similar if there is a utility-maximizing choice beginning with C at $(h,k+1)$. □

The contribution of Lemma 3.3 is to link the CS versus S comparison to backward induction. If either S is not an optimal choice at $(h+1,k)$ or S is not an optimal choice at $(h,k+1)$, then C is uniquely optimal at (h,k). But, for example, if $n = h + k - 1$, then S is mandatory at $(h+1,k)$ or $(h,k+1)$ since no more consultations are allowed, so both S and C are optimal choices at (h,k).

To study backward induction, it is convenient to separate cases according to the parity of n.

THEOREM 3.4. *If $c = 0$ and $n = 2m - 2$ consultations are permitted, where $m \geq 1$, DM's unique optimal choice is C at any state (h,k) satisfying* $\max\{h,k\} \leq m - 1$. *At every other state, both C and S are optimal.*

PROOF. If C is always chosen, then the final state (h,k) satisfies $h+k = 2m$. At state (m,m), C is not permitted, so the expected utility is $\frac{1}{2}$. By Lemma 3.2, if $1 \leq t < m$, DM's utility is increased by choosing C at state (t,t). By Lemma 3.3, the expected utility increases by choosing C at any state $(t-1,t)$, where $2 \leq t < m$. By induction, the expected utility increases by choosing C at any state (r,s), where $1 \leq s < r < m$ or $1 \leq r < s < m$. It is easy to see that such states (h,k) are exactly those that satisfy $\max\{h,k\} \leq m-1$. At all other states, both C and S are optimal, by Lemma 3.1. □

Note that DM is indifferent between S and C exactly when $\max\{h,k\} \geq m$.

Figure 4 shows DM's optimal strategy and expected utilities when $n = 4$. Since $m = 3$, Theorem 3.4 shows that the optimal choice is to consult at all states (h,k) where $\max\{h,k\} \leq 2$, that is, at (1, 1), (1, 2), (2, 1), and (2, 2). At all other states either C or S may be chosen. Each state is labelled with DM's expected utility at that state should it be attained when the optimal strategy is played. Notice that whether Consult is uniquely optimal depends on both the balance of evidence and the amount of evidence; the evidence favours Yes at (2, 1) and the unique optimal choice is Consult, whereas at (3, 2) the evidence still favours Yes, but now DM is indifferent between Consult and Stop.

THEOREM 3.5. *If $c = 0$ and $n = 2m - 1$ consultations allowed, where $m \geq 1$, DM's optimal strategy is C at any state (h,k) such that* $\max\{h,k\} \leq m$. *At every other state, both C and S are optimal.*

PROOF. If C is always chosen, then the final state (h,k) must satisfy $h+k = 2m+1$. At state (m,m), one more consultation is permitted. By Lemma 3.2, DM strictly prefers C at state (m,m), and at any state (t,t) where $1 \leq t \leq m$. By Lemma 3.3, the expected utility increases by choosing C at any state $(t-1,t)$, where $2 \leq t \leq m$. By induction, expected utility increases by choosing Consult at any state (r,s), where $1 \leq s < r \leq m$ or $1 \leq r < s \leq m$. At all other states, both C and S are optimal, by Lemma 3.1. □

Now it is the case that DM is indifferent between S and C exactly when $\max\{h,k\} > m$. Figure 5 shows DM's optimal strategy and expected utilities when $n = 4$. Since $m = 3$, Theorem 3.5 shows that the optimal choice is to consult at all states (h,k) where $\max\{h,k\} \leq 3$, that is, at (1, 1), (1, 2), (2, 1), (3, 1), (2, 2), and (1, 3). At all other states either C or S may be chosen.

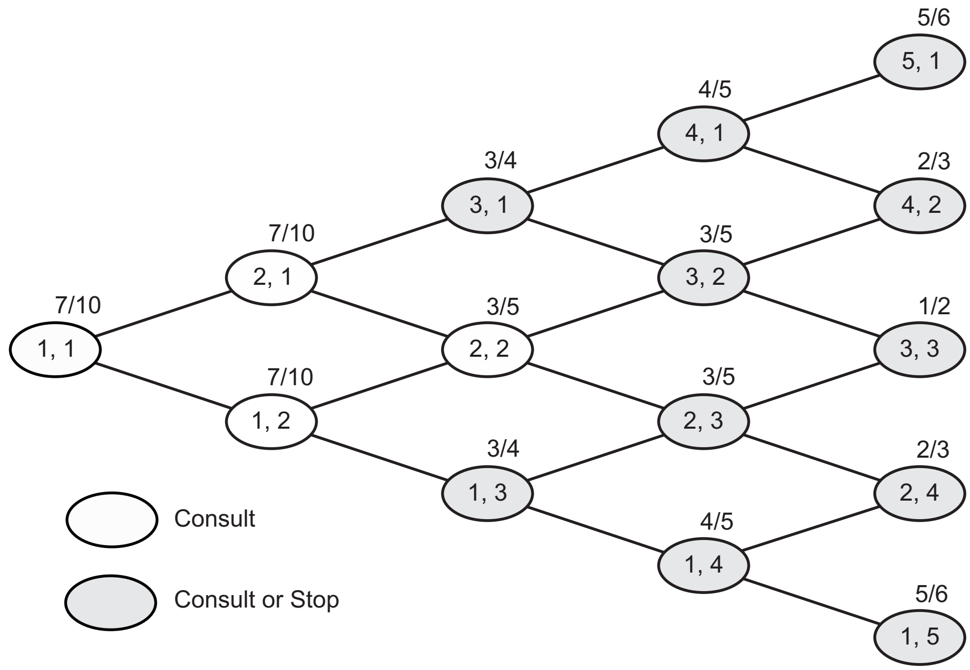

FIGURE 4. Optimal strategies and expected utilities when $n = 4$ and $c = 0$

To state the optimal strategy formally, define the set of *Consultation Nodes* to be

$$\mathcal{C}(n) = \left\{ (h,k) : h \geq 1, k \geq 1, \max\{h,k\} = \left\lfloor \frac{n+1}{2} \right\rfloor \right\}$$

For example, $(h,k) = (3,2) \in \mathcal{C}(5)$, because $\max\{h,k\} = 3 = \frac{5+1}{2}$, but $(h,k) = (4,2) \notin \mathcal{C}(5)$.

If $c = 0$ and the maximum number of consultations is n, a strategy is optimal if and only if it specifies C at all nodes of $\mathcal{C}(n)$. At all other nodes, either C or S may be chosen. Notice that, if n is odd, then $\mathcal{C}(n) = \mathcal{C}(n+1)$.

When n consultations are permitted and the cost per consultation is c, denote DM's expected utility using the optimal strategy by $U_c^*(n)$. We can now specify $U_0^*(n)$.

THEOREM 3.6.

$$U_0^*(n) = \begin{cases} \frac{3n+2}{4n+4} & \textit{if } n \textit{ even} \\ \frac{3n+5}{4n+8} & \textit{if } n \textit{ odd} \end{cases}$$

PROOF. The strategy to always Consult is optimal. By induction, after n consultations, each possible state $(n+1,1),(n,2),\ldots,(1,n+1)$ has probability $\frac{1}{n+1}$. Suppose that $n > 0$ is even, and set $n = 2m-2$ so $m \geq 1$. All possible final states, $(j, 2m-j)$ where $j = 1,2,\ldots,2m-1$, have probability $\frac{1}{2m-1}$. It follows

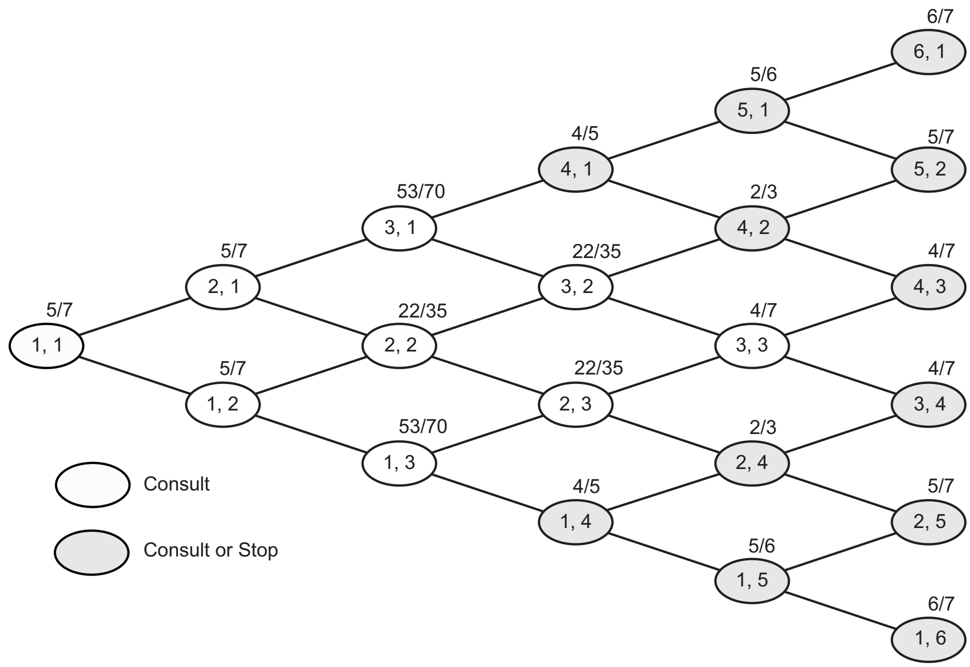

FIGURE 5. Optimal strategies and expected utilities when $n = 5$ and $c = 0$

that

$$U_0^*(n) = \frac{1}{2m-1} \sum_{j=1}^{2m-1} \frac{\max\{j, 2m-j\}}{2m} = \frac{3m-2}{4m-2}.$$

Substitution of $m = \frac{n+2}{2}$ now shows that $U_0^*(n) = \frac{3n+2}{4n+4}$ as required. Similarly, if n is odd, set $n = 2m - 1$ with $m \geq 1$. Then a similar calculation shows that

$$U_o^*(n) = \frac{1}{2m} \sum_{j=1}^{2m} \frac{\max\{j, 2m+1-j\}}{2m} = \frac{3m+1}{4m+2}$$

and substitution for $m = \frac{n+1}{2}$ now produces $U_0^*(n) = \frac{3n+5}{4n+8}$. □

Note that, if n is odd, so that $U_0^*(n) = \frac{3n+5}{4n+8}$, then $n+1$ is even, and $U_0^*(n+1) = \frac{3(n+1)+2}{4(n+1)+4} = \frac{3n+5}{4n+8} = U_0^*(n)$. Thus we have that, for example, $U_0^*(0) < U_0^*(1) = U_0^*(2) < U_0^*(3) = U_0^*(4) < U_0^*(5) \ldots$.

It is immediate from Theorem 3.6 that

$$U_0^*(n) \longrightarrow \frac{3}{4} \text{ as } n \longrightarrow \infty$$

For some intuition into this limit, recall that after n consultations, each possible state $(n+1, 1), (n, 2), \ldots, (1, n+1)$ has equal probability. In the limit, $\frac{h}{h+k}$ can be thought of as a uniform random variable Y on $[0, 1]$; the limiting value of $\frac{\max\{h,k\}}{h+k}$

is $M(X) = \max\{Y, 1-Y\}$. It is easy to check that the expected value of $M(Y)$ is $\frac{3}{4}$, confirming the limiting value obtained above.

As $U_0^*(n)$ is non-decreasing, it follows that, as more consultations are permitted, the expected utility (following the optimal strategy) increases toward an upper bound of $\frac{3}{4}$. No amount of consultation can possibly increase expected utility beyond this bound, making it a fundamental limit on the value of consultation in our model.

4. Cost Models

Now we turn to Cost Models, in which $c > 0$ and $n = 0, 1, \ldots$. Expected utility is denoted $U_c(n)$; if an optimal strategy is followed, DM's expected utility is denoted $U_c^*(n)$.

4.1. Cost Model: $n = 0$. If no consultations are permitted, there are no decisions, and expected utility is $U_c^*(0) = \frac{1}{2}$.

4.2. Cost Model: $n = 1$. If 1 consultation is permitted, the only choice is C or S at (1, 1). DM's expected utility then equals

$$\begin{cases} \frac{1}{2}, & \text{if S} \\ \frac{2}{3} - c, & \text{if C} \end{cases}$$

It is clear that the optimal strategy is C if $c \leq \frac{1}{6}$ and S if $c \geq \frac{1}{6}$, so that

$$U_c^*(1) = \begin{cases} \frac{2}{3} - c, & \text{if } c \leq \frac{1}{6} \\ \frac{1}{2}, & \text{if } c \geq \frac{1}{6} \end{cases}$$

Thus, if $c > 0$ is small enough, the optimal strategy is C, exactly the same as at $c = 0$, but if c is large enough, the optimal strategy changes from C to S. Note also that $U_c^*(1)$ is a continuous function of c.

We now compare the choices S and CS at a typical node (h, k).

LEMMA 4.1. *Assume $c > 0$ and $n > h + k - 2$. If $h \neq k$ then, at (h, k), S gives greater expected utility than CS.*

PROOF. At (h, k) there have been $h + k - 2$ consultations, so S yields expected utility $\frac{h}{h+k} - (h + k - 2)c$, whereas CS yields expected utility

$$\frac{h}{h+k}\left(\frac{h+1}{h+k+1} - (h+k-1)c\right) + \frac{k}{h+k}\left(\frac{h}{h+k+1} - (h+k-1)c\right)$$

which is easily shown to equal $\frac{h}{h+k} - (h+k-1)c$. The utility of S therefore exceeds the utility of CS by $c > 0$. The proof is analogous if $h < k$. $\square$

LEMMA 4.2. *Assume $n > 2h - 4$. At state (h, h), CS yields greater expected utility than S if and only if $c < \frac{1}{2(2h+1)}$.*

PROOF. The expected utility for S is $\frac{1}{2} - (2h - 2)c$, and the expected utility for CS is $\frac{h+1}{2h+1} - (2h-1)c$. It is easily shown that expected utility for CS is greater iff the specified condition is satisfied. $\square$

Thus, C is better at (h, h) provided c is small enough. In particular, our analysis of the $h = 1, n = 1$ case is confirmed, and it also applies to $n = 2$. Note that if it is optimal to choose C when $c > 0$, then it must be optimal to choose C when $c = 0$. Therefore, the nodes where it is optimal to choose C are always a subset of $\mathcal{C}(n)$.

Starting at a state (h, k) where $h \neq k$, let G be the conditional strategy of continuing to choose C until the state returns to the centre line, and then choosing C once more. (If $h > k$, proceeding directly to the centre line will produce the state (h, h); strategy G calls for one more choice of C when this state is attained.)

Lemma 4.3. *Assume $c > 0$ and $n > 2h-2$. At state $(h, h-1)$, G yields greater expected utility than S if and only if $c < \frac{h-1}{2(2h+1)(3h-2)}$.*

Proof. At state (h, k), the expected utility for S is

$$\frac{h}{2h-1} - (2h-3)c$$

The expected utility for G is

$$\frac{h}{2h-1}\left(\frac{h+1}{2h} - (2h-2)c\right) + \frac{h-1}{2h+1}\left(\frac{h+1}{2h+1} - (2h-1)c\right)$$

which can be shown to equal

$$\frac{4h-1}{2h-1}\left(\frac{h+1}{2(2h+1)} - (h-1)c\right).$$

After some manipulation, it can be seen that the expected utility for G is greater than for S iff

$$c < \frac{h-1}{2(2h+1)(3h-2)}$$

□

It follows from Lemmas 4.2 and 4.3 that the optimal strategy when $n = 3$ is given by

- At (1, 1), choose C if $c \leq \frac{1}{6}$ and S otherwise
- At (2, 1) and (1, 2), choose C if $c \leq \frac{1}{40}$ and S otherwise
- At (2, 2), choose C
- At all other nodes, choose S

The resulting utility is given by

$$U_c^*(3) = \begin{cases} \frac{7}{10} - 3c, & \text{if } c \leq \frac{1}{40} \\ \frac{2}{3} - c, & \text{if } \frac{1}{40} \leq c \leq \frac{1}{6} \\ \frac{1}{2}, & \text{if } \frac{1}{6} \leq c \end{cases}$$

Moreover, the same strategy is optimal when $n = 4$, and $U_c^*(4) = U_c^*(3)$. Again, note that $U_c^*(3)$ is a continuous function of c for $c \geq 0$.

This optimal strategy is consistent with the optimal choices shown (for $n = 3$) in Figure 6. Note that these choices are "subgame perfect" in the sense that they prescribe the optimal choice on arriving at each node. Should DM arrive at node (2, 2), the optimal choice is to consult as long as $c < \frac{1}{10}$. However, DM can reach (2, 2) only by consulting at (2, 1) or (1, 2), and at these two nodes the optimal choice is Consult only if $c < \frac{1}{40}$. Thus, if the optimal strategy (starting at (1, 1)) is followed, testing whether $c < \frac{1}{10}$ at (2, 2) is irrelevant; this node cannot be attained unless $c < \frac{1}{40}$. But if, perhaps as a result of error, (2, 2) is reached even though $c > \frac{1}{40}$, the optimal choice is Consult if $c < \frac{1}{10}$ and Stop if $c > \frac{1}{10}$.

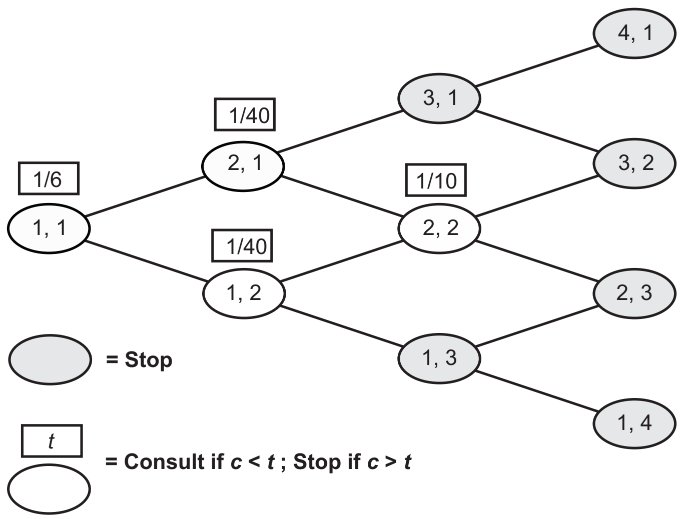

FIGURE 6. Optimal strategies and expected utilities when $n = 3$ and $c \geq 0$

When $n = 5$ and $c \geq 0$, a similar calculation yields Figure 7. In view of Figure 5, we anticipate that when $c \longrightarrow 0$, $U_c^*(5) \longrightarrow \frac{5}{7}$. This is the case, but the strategy to Consult at every node of $\mathcal{C}(5)$ is implemented only if $c < \frac{1}{189}$.

5. Conclusions

Our model suggests some general rules of thumb for decision-makers who may consult to gain evidence to support a decision.

- Consult when there is relatively little evidence, or when the evidence is approximately balanced
- Consult less as consultation becomes costlier, particularly when the evidence is less balanced

The deliberations over the U.S. response to the Soviet missiles in Cuba demonstrate that continued consultation can be important, and the choice to keep consulting or to act can be crucial. At the outset of the Cuban missile crisis, EXCOMM had 13 members, mostly cabinet secretaries, but it then expanded to include 13 so-called advisers, who were brought in to offer counsel on various matters. (Because these advisors came and went, they do not follow strictly our model that allows consultants only to be added.) The consensus that the committee reached was that if a blockade of Cuba failed to induce Soviet withdrawal of the missiles, it could be followed by more aggressive action—specifically, an air strike—to try to eliminate the threat that the missiles posed, but at the possible cost of provoking a nuclear war. After the United States secretly agreed to withdraw its comparable missiles from Turkey if the Soviet Union withdrew its missiles from Cuba, the crisis

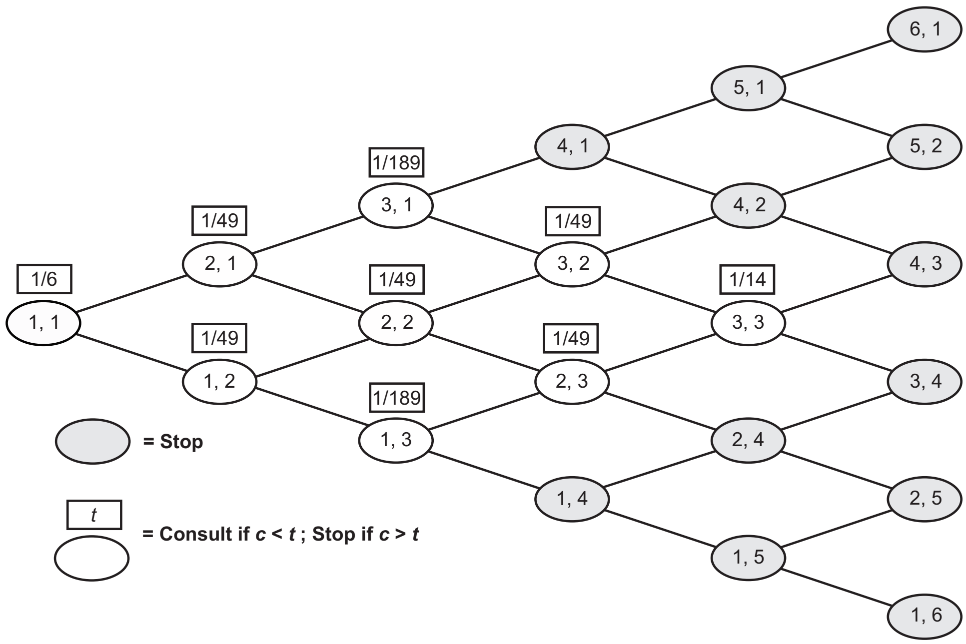

FIGURE 7. Optimal strategies and expected utilities when $n = 5$ and $c \geq 0$

subsided, illustrating how consultations, which occurred over 13 harrowing days, may lead to a propitious solution to a crisis.

One important conclusion from our model is that there is a fundamental upper limit on the utility of a decision-maker who is initially maximally uncertain (i.e., who begins at $(h, k) = (1, 1)$), no matter how much consultation is undertaken. This limit applies even if consultation is free ($c = 0$).

By being specific, our model makes these observations explicit. For example, when consultation is costless ($c = 0$) and n consultations are permitted, the rule of thumb indicates that a DM who begins at $(h, k) = (1, 1)$ should consult at any state (h, k) where

$$1 \leq h, k \leq \frac{n+1}{2}$$

To determine whether to consult at a state (h, k), a simple test is whether strategy G is feasible; if $h > k$ say, the question is whether, if there were an unbroken sequence of No's resulting in the state (h, h), would one more consultation be permitted? If the answer to this question is positive, then G is feasible, and—if c is low enough—the decision-maker should consult at (h, k).

Simply guessing initially gives a maximally uncertain decision-maker an expected utility of $\frac{1}{2}$. Consultation can raise this expected utility, but even with unlimited costless consultation it cannot exceed the fundamental limit of $\frac{3}{4}$. Consultation can increase expected utility if costs are zero, but it does so at a decreasing rate. Even following an optimal strategy, the upper bound on the expected utility approaches $\frac{3}{4}$, so perfection—always making the right decision—is unattainable through consultation because of the inherent uncertainties in the model.

We conclude with some comments on possible variations in our model, which we leave for a future project. First, a decision-maker who is not maximally uncertain may start at initial state (h_0, k_0) for any $h_0 > 0$ and $k_0 > 0$. Moreover, consultation may cause a state transition

$$(h,k) \longrightarrow \begin{cases} (h+\alpha, k) & \text{with probability } \frac{h}{h+k} \\ (h, k+\alpha) & \text{with probability } \frac{k}{h+k} \end{cases}$$

for some fixed $\alpha > 0$. To add more complexity, consultant ℓ may have cost c_ℓ and transition the state by an amount α_ℓ. But distinguishing consultants in this way may raise another decision problem, namely how DM should choose a consultant to get the most information, α_ℓ, at the least cost, c_ℓ.

We have assumed additive costs; another variation on our model could be multiplicative costs, so the net utility of Stop at (h, k) would be $u(X_{op})\cdot(1+\gamma)^{h+k-2}$ where $\gamma > 0$ is a new parameter. A hybrid model comprising both additive and multiplicative costs is also possible. We do not explore these variant of our model, but we conjecture that they have similar properties to the additive model.

Acknowledgement: The authors are grateful to the editors and an anonymous referee for comments that helped to improve the manuscript.

References

[A] Jeanne Pavy, *The Role of the Executive Committee in the Cuban Missile Crisis.* http://people.loyno.edu/ history/journal/1983-4/1983-4.htm/, accessed 2022-04-02.

[B] Robert F. Kennedy, *Thirteen Days*, W. W. Norton, 1969.

[C] Steven J. Brams, *Theory of Moves*, Cambridge University Press, 1993, pp. 130–38.

[D] F. Thomas Bruss, *A Unified Approach to a Class of Best Choice Problems with an Unknown Number of Options*, The Annals of Probability, 12, 3, 1984, pp. 882–89.

[E] Secretary Problem, Wikipedia, https://en.wikipedia.org/wiki/Secretary_problem, accessed 2022-04-28.

[F] Steven J. Brams and D. Marc Kilgour, Revisiting the Secretary Problem. Preprint, New York University.

[G] Beta-binomial Distribution, Wikipedia, https://en.wikipedia.org/wiki/Beta-binomial_distribution, accessed 2022-04-28.

Department of Mathematics, Wilfrid Laurier University, Waterloo, ON N2L 3C5, CANADA

E-mail address: mkilgour@wlu.ca

Department of Politics, New York University, New York, NY 10012, USA

E-mail address: steven.brams@nyu.edu