worthwhile for Scotia.

5. A simple bidding problem

Case: THE BATES RESTORATION (A)

Russ Gehrig, a construction general contractor, has decided to bid for the contract to do an extensive restoration of the old Bates mansion. The contract will be awarded to the bidder who submits the lowest bid, and the winner will be paid the amount of his bid at the completion of the job. Gehrig estimates that his cost to do the work described in the contract is equally likely to be above or below $100,000, has probability 1/4 of being below $86,000, and has probability 1/4 of being above $120,000.

Gerhig knows that there will be at least one other bidder for the job, but he does not know how many. He has assessed the following probability distribution for the number of other bidders (not including himself):

# Other Bidders	Probability
1	0.2
2	0.3
3	0.3
4	0.1
5	0.1

He also figures the each opposing bidder's bid is equally likely to be above or below $140,000, has probability 1/4 of being below $120,000, and has probability 1/4 of being above $180,000.

To assess his risk tolerance, Gerhig thought about a simple hypothetical gamble that would either generate a loss of $20,000, or generate some profit of $X for him, each with probability 1/2. He tried to ask himself, what was the smallest number X such that he would be willing to accept such a gamble? If the potential profit $X was less than $20,00, then the expected monetary value of the gamble would be less than $0, and he would not accept such a gamble. Indeed, he might still walk away from the gamble even if $X was slightly larger than $20,000. But if the potential profit $X was $45,000 or more, then he would feel willing to accept such a gamble.

Selecting a bid for an auction is a decision under uncertainty. A bidder might think about submitting an aggressive bid that would have a high probability of winning the auction, but such an aggressive bid might yield the bidder little or no profit in the event that it wins. On the other hand, a less aggressive bid that would allow a larger profit margin for the bidder typically would have a smaller probability of winning. So our criterion for choosing a bid must take account of both the probability of winning and the potential profit that the bid would yield if it wins.

One such criterion for determining optimal bids is expected profit maximization. In an auction where losing bidders pay nothing and get nothing, a bidder's expected profit from submitting some bid β is the probability of the bid β winning the auction multiplied by the conditionally expected profit that the bidder would get if he won the auction with this bid β.

To be more precise, let us consider the auction described in part (A) of the "Bates Restoration" case. The auction in this case is an example of a first-price sealed-bid auction in which the bidders are bidding to sell. In an auction where the bidders are bidding to sell, the winner of the auction will be the one with the lowest bid. In a sealed-bid auction, the bids are submitted independently, and there is no opportunity to revise one's bids after hearing others' bids. In a first-price auction where the bidders are sellers, the winner will be paid the amount of his bid to provide the services specified in the contract. Of course, the bidders who do not win are paid nothing and have no obligations to provide any services.

So suppose that we are consultants for the bidder (Gehrig) in the Bates case. Let β denote his bid, let **C** denote his cost of fulfilling the contract, and let **A** denote the lowest opposing bid submitted by other bidders. The bid β is a decision variable for Gehrig, but the cost **C** and lowest opposing bid **A** must be treated as random variables (as conventionally indicated here by boldface) because our decision-maker does not know them. With this notation, the profit from submitting the bid β in this auction is

$$\text{Profit} = \beta - \mathbf{C} \ \text{ if } \ \beta < \mathbf{A}, \ \text{ and } \ \text{Profit} = 0 \ \text{ if } \ \beta > \mathbf{A}.$$

For a risk-neutral bidder, the optimal bid is one that maximizes the bidder's expected profit from submitting the bid β, which is

$$\text{E(Profit from submitting bid } \beta) = \text{P}(\mathbf{A}>\beta) * (\beta - \text{E}(\mathbf{C} \mid \mathbf{A}>\beta)).$$

This formula expresses a natural trade-off between the goals of having a high probability of winning, and of having a large profit margin when the bid wins. Choosing a lower bid β could increase the probability of winning $\text{P}(\mathbf{A}>\beta)$, but it also would decrease the profit $(\beta-\text{E}(\mathbf{C} \mid \mathbf{A}>\beta))$

that he could expect when he wins.

Figure 7 shows an analysis of this bidding problem. The cost of fulfilling the contract is simulated in cell A1 of Figure 7 by a Generalized-Lognormal random variable with Gehrig's assessed quartiles for the cost (86, 100, and 120 thousand dollars). The uncertain number of bidders is simulated in cell A8, using the discrete probability distribution shown in cells C10:D14. The opposing bids are simulated in cells A10:A14 of Figure 7 by entering the formula

=IF(C10<=$A$8, GENLINV(RAND(),$F$10,$F$11,$F$12), "..")

into cell A10, and then copying cell A10 to A10:A14. Notice that cells C10:C14 contain the numbers 1 to 5. So given any number from 1 to 5 in cell A8, the cells A10:A14 will contain as many simulated bids as this number in cell A8 (with each bid being an independent Generalized-Lognormal random variable with the assessed quartile points 120, 140, 180); and the (5−A5) extra cells in A10:A14 will contain a nonnumerical text (".."). The formula =MIN(A10:A14) in cell A16 then returns the lowest opposing bid in this simulation model. (The MIN function simply ignores any nonnumerical cells in the A10:A14 range.)

A simulation table in A21:C2022 contains 2001 simulated values of the potential cost (in B22:B2022) and the lowest opposing bid (in C22:C2022). A proposed bid has been entered into cell D18 of Figure 7. For each row of simulation data, the question of whether the D18 bid would win the auction in this simulation is answered (1=yes, 0=no) in column E by entering

=IF(C22>$D$18,1,0)

in cell E22, and then copying E22 to E22:E2022. The corresponding profit for Gehrig in each simulation row is computed in column F by entering the formula

=E22*($D$18−B22)

into cell F22, and then copying F22 to F22:F422. This formula in F22 returns $0 whenever Gerhig's D18 bid would lose the auction (when E22=0), and returns Gerhig's net profit from the contract (D18−B22) whenever his bid would win (when E22=1). So Gerhig's expected profit from the bid in cell D18 can be estimated in cell I34 by the formula

=AVERAGE(F22:F2022)

	A	B	C	D	E	F	G	H	I	J	K	L	M
1	95.91624	Gerhig's Cost to do the Bates job ($1000s)											
2					Quartiles for Gerhig's actual cost								
3					86								
4					100								
5					120								
6													
7													
8	3	#Other bidders											
9	Other bids		#Bidders	Proby		Quartiles for each other bid							
10	159.9106		1	0.2		120							
11	132.4324		2	0.3		140							
12	137.2995		3	0.3		180							
13	..		4	0.1									
14	..		5	0.1				FORMULAS FROM RANGE A1:G22					
15								A1. =GENLINV(RAND(),E3,E4,E5)					
16	132.4324	Lowest opposing bid						A8. =DISCRINV(RAND(),C10:C14,D10:D14)					
17								A10. =IF(C10<=$A$8,GENLINV(RAND(),$F$10,$F$11,$F$12),"..")					
18				140	G's bid			A10 copied to A10:A14.					
19								A16. =MIN(A10:A14)					
20		G's cost	Bid to beat		G wins?	G's profit		B21. =A1					
21	SimTable	95.91624	132.4324					C21. =A16					
22	0	83.59129	120.6856		0	0		E22. =IF(C22>$D$18,1,0)					
23	0.0005	106.3137	116.1297		0	0		F22. =E22*($D$18-B22)					
24	0.001	94.6027	201.7363		1	45.3973		E22:F22 copied down.					
25	0.0015	89.33786	126.9999		0	0							
26	0.002	98.71699	138.3054		0	0							
27	0.0025	119.5908	156.9625		1	20.40922							
28	0.003	148.5063	122.3162		0	0							
29	0.0035	94.14449	118.6837		0	0					FORMULAS FROM H30:J34		
30	0.004	89.92413	114.0844		0	0		RiskTol	CE from bid		H31. =RISKTOL(45,-20,0)		
31	0.0045	212.6046	123.1261		0	0		38	2.346695		I31. =CE(F22:F2022,H31)		
32	0.005	71.09629	112.1275		0	0					H34. =AVERAGE(E22:E2022)		
33	0.0055	85.41878	110.1573		0	0		Pr(Win)	E Profit	Stdev	I34. =AVERAGE(F22:F2022)		
34	0.006	106.755	134.255		0	0		0.233383	7.429879	20.06721	J34. =STDEV(F22:F2022)		

Figure 7. Bidding model for the Bates case, part A (no winner's curse).

Gerhig has indicated that a gamble that would either generate a profit of about $45 thousand or a loss of $20 thousand, each with probability 1/2, might not be substantially better than a sure alternative of $0. So we can estimate that his risk tolerance is approximately $38 thousand, as computed by the formula =RISKTOL(45,−20,0) in cell H31 of Figure 7. Then Gerhig's certainty equivalent for the gamble that he gets by submitting the bid in cell D18 can be computed in cell I31 by the formula

=CE(F22:F2022,H31)

The expected profit in cell I34 and the certainty equivalent in cell I31 in Figure 7 have been recalculated for a range of bids from $100 thousand to $300 thousand, with a fixed table simulation data in B22:C2022. The results are shown in Figure 8. If Gerhig's goal is to maximize expected profit, then a bid of $140 thousand seems to be optimal, yielding an expected profit of about $7.43 thousand. On the other hand, if Gerhig's limited risk tolerance is taken into account, then we find that a somewhat higher bid may be optimal. With this simulation, we find a maximal CE of $2.83 thousand, when the bid is $155 thousand.

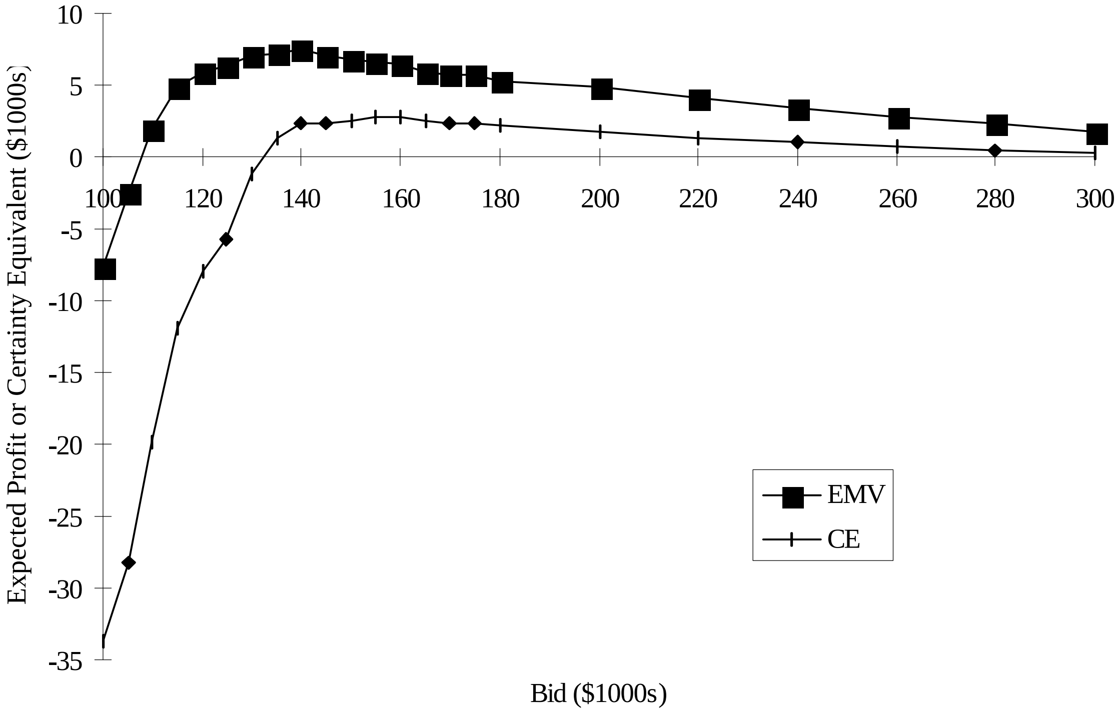

Figure 8. Finding optimal bids for the Bates case, part A (no winner's curse).

6. The winner's curse

In the previous section we considered an auction where bidders were bidding to sell. Now let us consider the problem of someone who is trying to buy something in a bargaining situation, taking a well-known example from an article on "The Market for Lemons" by George Akerlof (1970). In this example, we are trying to help someone who has been negotiating to buy a used car. The problem is that the buyer has substantial uncertainty about the condition of this used car, and the buyer recognizes that the seller has much better information about the car's condition, because the seller has owned it for a long time. (A used car of poor quality is often called a "lemon" in colloquial American English.)

Bargaining processes can be quite complicated, so let us simplify the situation by supposing that our buyer can make a firm take-it-or-leave-it offer to buy at some price. Then the buyer's only question is what price to offer the seller, and this price is like a bid in an auction. If the seller understands that this price is the buyer's final offer, then the seller should accept it if the offered price is more than the seller's own value (or reserve price) for keeping the car. If this were an auction, then our buyer would win the car if her bid were higher than the highest opposing bid. (In a situation of bidding to buy, it is the highest bid that wins!) So we can view this bargaining situation as a simple auction in which this final offer price is the buyer's bid, and the seller's value for keeping the car is the opposing bid that our buyer's bid must exceed to get the car. If the buyer's bid is greater than the seller's value, then the buyer's profit is the buyer's value of the car minus the buyer's bid. But if the buyer's bid is less than the seller's value then the buyer gets zero profit because there is no transaction.

As long as the car is worth more to the buyer than the seller, the buyer's ideal bid would be one cent more than the seller's value for keeping the car. But of course the buyer does not know the seller's reserve price. So this bargaining problem is a decision under uncertainty for the buyer. To find an optimal bid that maximizes the buyer's expected profit (assuming risk neutrality), we must ask the buyer to assess probability distributions for the seller's reserve price and the buyer's value of getting the car.

Suppose that the (female) buyer tells us that the value of the car to her will depend on the quality of its condition (which determines how long it will run before breaking down). Because she does not know the actual condition of the car, she thinks that it could be worth anywhere from \$0 to \$1500 to her. The buyer also believes that the (male) seller's value of keeping the could be anywhere from \$0 to \$1000, also depending on the condition of the car (which the seller knows). In each case, the buyer is willing to assess a Uniform distribution for these two unknowns. Let **V** denote the value that the car could yield for the buyer, and let **A** denote the seller's reserve value for keeping the car. With this notation, the buyer's profit from a bid β is

$$\text{Profit} = \mathbf{V} - \beta \text{ if } \beta > \mathbf{A}, \text{ and Profit} = 0 \text{ if } \beta < \mathbf{A}.$$

Because the buyer's value **V** is drawn from a Uniform distribution from \$0 to \$1500, the expected value of the car to the buyer is

$$E(\mathbf{V}) = 750.$$

For any bid β that is between \$0 and \$1000, the probability of β being greater than the seller's reserve value **A** is

$$P(\mathbf{A}<\beta) = \beta/1000$$

because **A** is drawn from a Uniform distribution from \$0 to \$1000.

Consider now the consequences of bidding $\beta = 375$. This bid β happens to be the value that maximizes the quantity $(E(\mathbf{V})-\beta)*P(\mathbf{A}<\beta)$, and it would be accepted by the seller with probability 375/1000. So the buyer has probability 0.375 of being able to buy the car for \$375. Since the expected value of the car for her is \$750, we might initially guess that this bid should give her a positive expected profit.

There is a serious problem with this analysis, however, because we have not assessed the relationship between the seller's value **A** and the buyer's value **V**. Both are unknown quantities in the buyer's model because they depend on the condition of the car, which is known to the seller but not to the buyer. So there is good reason to suspect that these two random variables should be highly correlated. To make things simple, suppose that the buyer tells us that **A** and **V** have a perfect correlation of 1, because the value of the car to her would always be 50% more than the value of the car to the seller. That is, suppose that **V** depends on **A** by the simple linear relation

$$\mathbf{V} = 1.5 * \mathbf{A}$$

With this equation, having **A** drawn from a Uniform distribution over \$0 to \$1000 does indeed make **V** a random variable with a Uniform distribution between $0 and $1500. Notice also that this equation implies that the car would always have more value to the buyer than to the seller.

But a simulation analysis of this bargaining situation, shown in Figure 9, indicates that the buyer's expected profit from bidding $375 is actually negative. An average of the results of 401 simulations, computed in cell E10 of Figure 9, indicates that bidding $375 would actually yield an expected net loss of about $35!

	A	B	C	D	E	F
1	AKERLOF'S LEMON PROBLEM					
2	FORMULAS					
3	B15. =1000*RAND()					
4	C15. =B15*1.5					
5	E16. =IF($D$15>B16,C16-$D$15,0)				E(ValueToB|BidWins)	
6	F16. =IF($D$15>B16,C16,"..")					283.53
7	E16:F16 copied to E16:F416					P(BidWins)
8	E10. =AVERAGE(E16:E416)					0.3791
9	F6. =AVERAGE(F16:F416)				E(BsProfit)	
10	F8. =COUNT(F16:F416)/COUNT(C16:C416)				-34.67	
11	E11. =F8*(F6-D15)				-34.67	
12						
13		Seller's value				
14		(reserve price)	Value to Buyer	Buyer's bid		
15	SimTable	441.94	662.91	375	B's Profit	ValToB|Win
16	0	806.04	1209.06		0.00	..
17	0.0025	550.15	825.23		0.00	..
18	0.005	242.42	363.64		-11.36	363.64
19	0.0075	810.76	1216.14		0.00	..
20	0.01	429.47	644.20		0.00	..
21	0.0125	870.94	1306.41		0.00	..
22	0.015	634.63	951.94		0.00	..
23	0.0175	865.50	1298.25		0.00	..
24	0.02	108.56	162.85		-212.15	162.85
25	0.0225	666.59	999.89		0.00	..
26	0.025	534.08	801.11		0.00	..
27	0.0275	166.38	249.57		-125.43	249.57
28	0.03	215.52	323.28		-51.72	323.28
29	0.0325	374.35	561.53		186.53	561.53
30	0.035	100.58	150.87		-224.13	150.87
31	0.0375	512.42	768.64		0.00	..
32	0.04	896.87	1345.30		0.00	..

Figure 9. Bidding for a lemon.

What has gone wrong? The problem is that our bid is accepted only when the seller's value **A** for the car is relatively low (less than \$375), in which case the buyer's value **V** must also be relatively low (less than \$562.50 = 1.5*375), and the mean of these relatively low values is substantially less than \$750. The conditional expectation $E(\mathbf{V}|\mathbf{A}<375)$ has been estimated from simulation data in cell F9 of Figure 9, and these results show that, when the bid \$375 is accepted, the conditionally expected value of the car for the buyer is about \$280. So learning that the bid of \$375 is accepted should make the buyer expect to lose almost \$100 from the transaction!

The general formula for the buyer's profit from submitting a bid β is

$$E(\text{Profit}) = P(\mathbf{A}<\beta) * (E(\mathbf{V}|\mathbf{A}<\beta) - \beta).$$

This formula is applied in cell E11 of Figure 9, using the simulation estimates of $P(\mathbf{A}<\beta)$ and $E(\mathbf{V}|\mathbf{A}<\beta)$ from cells F8 and F6, and the result in cell E11 is indeed the same as the expected profit that we estimated in cell E10 using the same simulation data. When these calculations are repeated with any positive bid β entered into cell D15 of Figure 9, the expected buyer's profit in cell E10 and E11 is always negative. So the optimal bid for the buyer in this example is \$0.

The problem with our initial guess was that we used $E(\mathbf{V})$ instead of $E(\mathbf{V}|\mathbf{A}<\beta)$ in the formula for expected profit, and $E(\mathbf{V}|\mathbf{A}<\beta)$ is typically less than $E(\mathbf{V})$. That is, learning that a bid β has been accepted by the seller may make the buyer revise downward her expected value of the object that she is buying. This inequality

$$E(\mathbf{V}|\mathbf{A}<\beta) < E(\mathbf{V})$$

is called the winner's curse. Such a winner's curse can arise in bidding problems whenever the value of winning to the bidder is positively correlated with the other opposing bids. In this example, if the buyer's value **V** had been independent of the seller's value **A**, then there would have been no winner's curse, because learning that **A** was less than the bid β would convey no new information about the value **V** for the buyer.

Now let us reconsider the Bates case. Notice that the cost in cell A1 of Figure 7 depends on its own RAND, which is independent of the RANDs that drive the bids in A10:A14 of Figure 7. So this model implicitly assumes that Gerhig's cost of doing the Bates job would be independent of the opposing bids, which means that we have implicitly assumed that there would be no winner's curse in this case. Given the importance of the winner's curse in Akerlof's "lemon" example, this independence assumption should be seriously scrutinized.

Recall that two random variables are said to be independent if learning about one of them would not change our beliefs about the other. Should Gerhig feel this way about the his cost of the Bates job and the various opposing bids that are submitted by other contractors? There are very good reasons to think that he might not consider these unknown quantities to be independent. After all, the opposing bids will be generated by other contractors using their own best estimates of the costs of doing the Bates job, and there is some chance that these other contractors may have noticed some difficulties in the job which Gerhig has overlooked. Thus, if all the opposing bids turn out to be surprisingly high, then Gerhig might infer that the other contractors have generated higher cost estimates for the job, which might make him feel that his own cost estimates were probably too low.

To assess a possible winner's curse effect in this case, we could ask the decision-maker to think about how he might revise the median value of his cost distribution if he learned about some relatively high opposing bid. For example, we might ask Gerhig the following question:

> "Imagine that you could see the bid of one other contractor for the Bates job, and this contractor's bid turned out to be $180 thousand (which was the 0.75-percentile point of your assessed distribution for these bids). After getting this information, would you still feel that your own cost of doing the job was equally likely to be above or below $100 thousand (which was the 0.50-percentile point in the originally assessed distribution for this cost)? If not, then for what number X would you say that your cost was equally likely to be above or below $X, given this information about another bid of $180 thousand?"

Suppose that, in answer to these questions, Gerhig tells us that the conditional median of his cost to do the Bates job would increase to X = $105 thousand, given the information that another contractor was submitting a bid of $180 thousand. The fact that Gerhig's beliefs about his costs would be affected by information about opposing bids is an indication that the independence assumption in Figure 7 was wrong, and the model needs to be revised to take account of the winner's curse. Such a revised model is shown in Figure 10.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1	0.7737614	rand		122.76	Gerhig's Cost to do the Bates job ($1000s)									
2								Quartiles for Gerhig's actual cost						
3	0.2854	Correlation(Cost,OtherBid)						86						
4	Correl assessed from "0.5 = P(Cost<105|OtherBid=180)"							100						
5		Given value	Conditional median					120						
6	Corand	0.75	0.5763	Corand										
7	Other bid	180	105	Cost										
8				2	#Other bidders									
9	corands			Other bids		#Bidders	Proby		Quartiles for each other bid					
10	0.1613152			114.48		1	0.2		120					
11	0.8092553			198.32		2	0.3		140					
12	0.7531146			..		3	0.3		180					
13	0.445499			..		4	0.1							
14	0.6903316			..		5	0.1	FORMULAS FROM RANGE A1:I23						
15								A1. =RAND()						
16				114.48	Lowest opposing bid			D1. =GENLINV(A1,$H$3,$H$4,$H$5)						
17								C6. =MIDRAND(A3,B6) or =NORMSDIST(A3*NORMSINV(B6))						
18				140	G's bid			B7. =GENLINV(B6,I10,I11,I12)						
19								C7. =GENLINV(C6,H3,H4,H5)						
20		G's cost	Bid to beat		G wins?	G's profit		D8. =DISCRINV(RAND(),F10:F14,G10:G14)						
21	SimTable	122.763	114.48					A10. =CORAND($A$3,$A$1)						
22	0	130.339	141.29		1	9.6615		D10. =IF(F10<=$D$8,GENLINV(A10,$I$10,$I$11,$I$12),"..")						
23	0.0005	101.66	126.2		0	0		A10:D10 copied to A10:D14.						
24	0.001	171.719	139.89		0	0		D16. =MIN(D10:D14)						
25	0.0015	126.43	165.97		1	13.57		B21. =D1		C21. =D16				
26	0.002	100.077	111.94		0	0		E22. =IF(C22>$D$18,1,0)						
27	0.0025	87.0905	123.35		0	0		F22. =E22*($D$18-B22)						
28	0.003	90.1703	104.31		0	0		E22:F22 copied down.						
29	0.0035	120.154	128.33		0	0					FORMULAS FROM RANGE H30:J34			
30	0.004	107.662	142.95		1	32.338		RiskTol	CE from bid		H31. =RISKTOL(45,-20,0)			
31	0.0045	103.12	107.22		0	0		38	1.333774		I31. =CE(F22:F2022,H31)			
32	0.005	137.89	134.73		0	0					H34. =AVERAGE(E22:E2022)			
33	0.0055	113.827	110.84		0	0		Pr(Win)	E Profit	Stdev	I34. =AVERAGE(F22:F2022)			
34	0.006	94.2616	108.39		0	0		0.234414	5.101627	17.47748	J34. =STDEV(F22:F2022)			

Figure 10. Bidding model for the Bates case (part B) with a winner's curse effect.

The first problem is to find what correlation would correspond to the conditional median assessment that Gerhig has given us. The appropriate correlation has been entered into cell A3 of Figure 10, and its value has been derived from Gerhig's conditional-median assessment using some calculations shown in the A4:G7 box in Figure 10. To understand this box, notice first that cells B7 and C7 contain respectively the formulas

=GENLINV(B6,I10,I11,I12) and =GENLINV(C6,H3,H4,H5)

Cells I10:I12 contain Gerhig's assessed quartile points for his opposing bids, and cells H3:H5 contain Gerhig's (originally) assessed quartile points for his cost of doing the Bates job. So cells B7 and C7 would be random variables with the same distribution as Gerhig's cost and an opposing bid if B6 and C6 contained RANDs or CORANDs.

Now suppose that we revised the spreadsheet in Figure 10 by entering the array formula

{=CORAND(A3)}

into cells B6:C6. So cells B7 and C7 would now contain simulated bid and cost values that have the (normalized-rank) correlation in cell A3. To match Gerhig's conditional median assessment in this model, the appropriate value of A3 should be the correlation such that, when the CORAND in B6 makes B7 equal to 180, then the CORAND in C6 would make C7 equally likely to be greater or less than 105. We know that B7 is equal to 180 when B6 is equal to 0.75, because 180 is the 0.75-percentile point of the bid distribution. For such a pair of CORANDs in B6:C6 that have correlation A3, if one CORAND is given to have the value 0.75 then the other CORAND is equally likely to be greater or less than the value

=NORMSDIST(A3*NORMSINV(0.75))

which can also be returned (in Simtools v3.2) by the formula

=MIDRAND(A3, 0.75).

Then the value of A3 that we want is one such that GENLINV(MIDRAND(A3,0.75),H3,H4,H5) is equal to 105.

In Figure 10, the CORANDs in B6:C6 have been replaced by the value 0.75 in B6 and the formula =MIDRAND(A3,B6) in cell C6. Then the correlation in cell A3 has been adjusted by trial and error until a value close to 105 appears in cell C7. In this way, we have found that a correlation of 0.2854 generates random variables for bid and cost that have the relationship that

Gerhig described in his answer to our conditional-median question. Higher correlations in cell A3 would make the conditional median in cell C7 larger than 105, and lower correlations in A3 would make the conditional median in C7 less than 105.

So we want each of the opposing bids in our model to have correlation 0.2854 with the cost variable. But what correlation should the bids have with each other? Having higher correlation among the opposing bids would make them effectively act like a smaller number of bidders. Fewer opposing bidders would be good news for Gerhig, but perhaps he not be encouraged to such wishful thinking. So for simplicity and to be conservative, let us assume that the other bids are conditionally independent of each other given the cost variable. That is, let us assume that the other bids are only related by their shared correlation with the underlying cost of the job.

CORAND has an optional second parameter, called RandSource, which can be used to generate several CORANDs that are all correlated with one RAND but have no other relationship with each other. This formulation is used in cells A10:A14 in Figure 10. In Figure 10, Gerhig's cost of doing the Bates job is simulated in cell D9 by a GENLINV formula that is driven by a RAND in cell A1. Then the formula

```
=CORAND($A$3,$A$1)
```

has been entered into each of the cells in A10:A14. In this CORAND formula, cell A3 is the correlation parameter, and cell A1 is the optional RandSource parameter. Then CORAND(A3,A1) returns a Uniform random variable between 0 and 1 that is correlated with the RAND in cell A1 according to the correlation in cell A3.

These CORANDs in A10:A14 are then used to drive the simulated opposing bids in cells D10:D14 of Figure 10. Thus, with the value 0.2854 in cell A3, the opposing bids in cells D10:D14 become random variable that each have normalized-rank correlation 0.2854 with Gerhig's cost in cell D1. In this model, the opposing bids are not independent, because they all are all dependent on the same random variable that drives Gerhig's cost. But the bids in this model are conditionally independent given the cost, because they are driven by different CORANDs that are only linked by their common use of A1 as the secondary RandSource.

The rest of the model in Figure 10 is essentially the same as the model in Figure 7. The results of an analysis of 2001 simulations of this revised model are shown in Figures 10 and 11. Comparing Figure 11 to Figure 8, we can see that the winner's curse effect in this case tends to decrease Gerhig's expected profits. The expected profit is again maximized by a bid of $140 thousand, but the resulting maximal expected profit is now $5.10 thousand in Figure 10 (compared to a maximal expected profit of $7.43 thousand in the model without the winner's curse in Figure 7). When Gerhig's limited risk tolerance is taken into account, the results charted in Figure 11 suggest that Gerhig should can get a maximal CE of $2.35 thousand by using an optimal bid of about $150 thousand. So with either criterion (EMV or CE maximization), the optimal bid for this example seems to be relatively insensitive to the winner's curse effect.

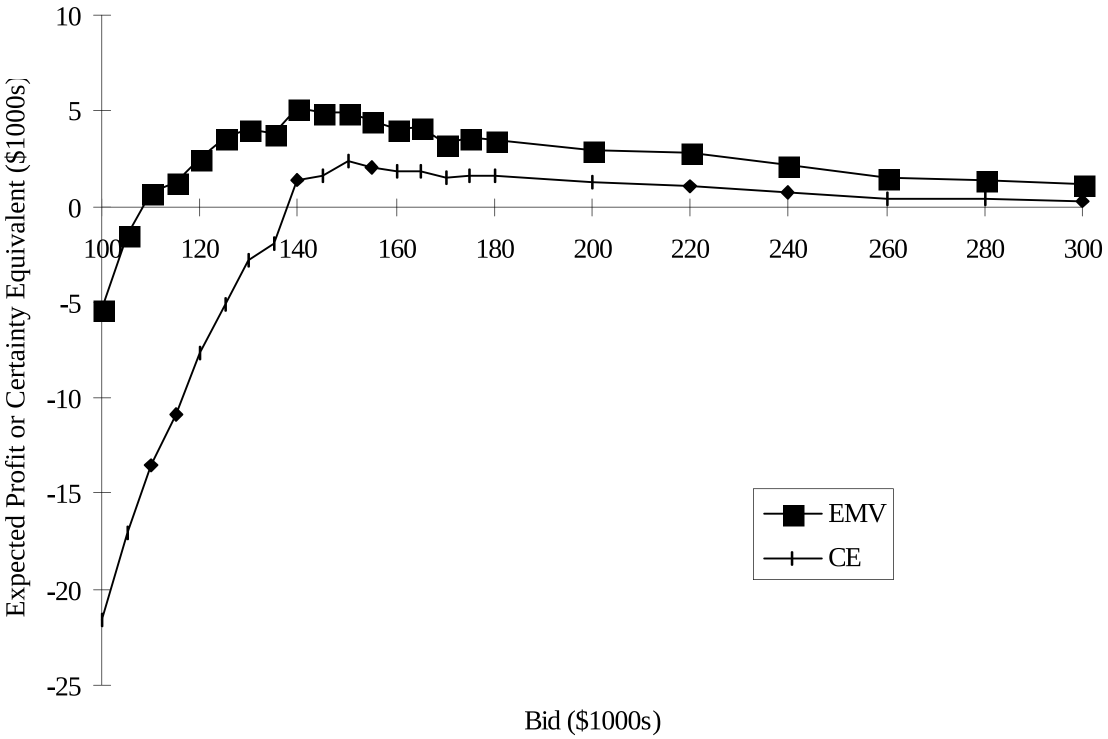

Figure 11. Finding optimal bids for the Bates case (B) with winner's curse.