The Winner's Curse

By Christian Svendsgaard

What kind of marketplace leaves the winner paying too much? It's called an auction, and the principle operates in the world of insurance as well as estates.

After my mother died, my sister sold her house at auction. She got seven bids, ranging from $625,000 to $757,000, with an average of about $695,000. All the bidders were bidding on the same house. They had the same information available to them. Yet their bids varied.

If we assume that the house has the same value to each bidder and that the average bid is the best estimate of the value of the house, then the winning bidder paid more than $60,000 too much. This difference between the mean bid and the winning bid is an example of the "winner's curse."

After he appeared on the cover of *SI*, Clint Mathis only started one game in the 2002 World Cup.

The theory of the winner's curse has been applied to many different situations: explaining the generally disappointing returns on leases of plots of land in the Gulf of Mexico for oil exploration (the leases are sold at auction); explaining why capital spending projects tend to have returns below expectations; explaining why the more job candidates you interview, the better your final choice (but also the more likely that that candidate won't live up to your expectations), etc. You could argue that the so-called "*Sports Illustrated* Cover Jinx"—that athletes have a disappointing game/match/race right after they're featured on the cover of *Sports Illustrated* magazine—is an example of the winner's curse.

The winner's curse also applies to insurance (and reinsurance). When we price a risk, we make our best estimate of what the losses will be. But various competitors will make different estimates. When the lowest price is the one that gets sold, the winner's curse is operating. The form of the auction is reversed from the house example, but otherwise there's little difference.

The winner's curse is studied in a branch of economics called auction theory.

The basic idea of auction theory is to treat the bids as random variables. Just as the losses insured under a contract are random, so too are the premiums.

The source of the randomness in the losses is more obvious, though. For setting insurance premiums—the selection of which underlying data to use, the model(s) fitted to it, the method used to allocate expenses, the decision about how much profit to seek, and the role of judgment—all these lead to the bids varying randomly. (Although it's certainly easier to see the randomness in someone else's bid than in one's own.) And because there's randomness in the bidding process, the winner's curse will apply.

Auction theory—the theory of the winner's curse—s a powerful analytical tool. Many applications can be made to insurance. Auction theory gives insight into the value of accuracy in pricing. It helps explain risk-averseness, and it illuminates the value of using actuarial credibility procedures and doing premium comparisons.

Is it Worth Hiring Actuaries?

Analyzing the winner's curse is a way of putting a value on the work of actuaries. At Swiss Re, we have a large and active facultative reinsurance department. Facultative reinsurance is written on individual risks, in contrast to treaties that cover portfolios of risks. The reinsurance contracts that cover facultative risks are called certificates—"fac certs" for short. Although there are some huge fac certs, most are much smaller than typical treaties. Because the contracts are so much smaller, there's little money available for internal expenses.

The question arises: Should an actuary help price a particular kind of fac cert? Traditional ratemaking science isn't useful in making this decision. It's clear that there's a cost to having an actuary work on a fac cert. But what's the benefit? The traditional attitude seems to be that more accuracy is always desirable, but

and Insurance

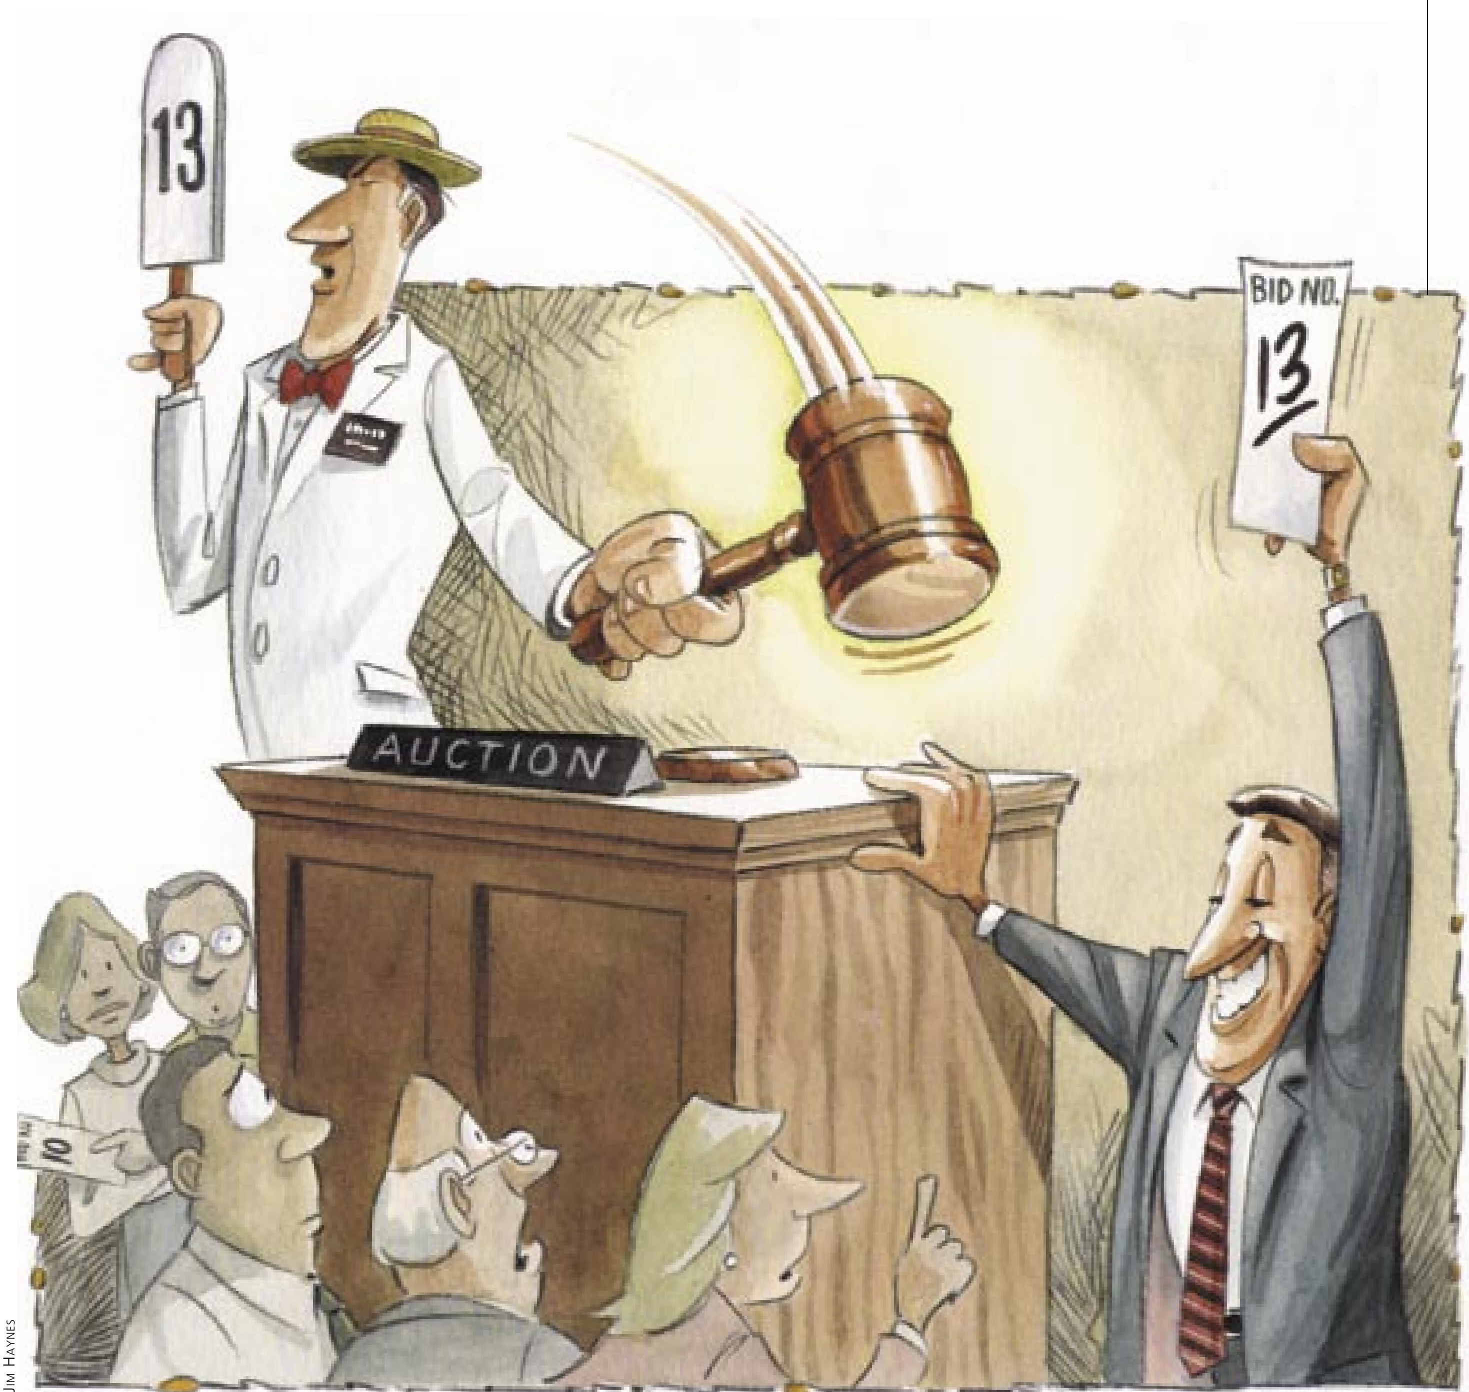

Jim Haynes

doesn't say how much that's worth in dollars. Analyzing the winner's curse can help put a price tag on increased accuracy.

Imagine that five reinsurers are bidding on a fac cert. The lowest bidder will win the entire contract. If all five reinsurers are essentially equivalent in their ability to price accurately—if their bids are independently and identically distributed—then each bidder has a 20 percent chance of winning the contract. The winning bid will be the minimum of the five bids, so the mean winning bid will be less than the average bid, and the difference is the winner's curse.

An extreme value from a larger population will tend to be more extreme. Yao, from China (pop. 1.2 billion), is 7'5"; Ewing, from Jamaica (pop. 2.7 million), is "only" 7' 0".

Now assume that one of the five competitors hires an actuary, so that the variance of that competitor's bid is cut in half. What happens?

All the bidders need to adjust their bidding strategies. If they don't, two things will happen. First, the smarter competitor will suffer less bias; second, the smarter competitor will be the winning bidder less often, because he won't screw up and give out too low a bid as often.

It's not as clear what happens if the bidders adjust their bidding strategies (although the chances of the non-actuary-employing bidders correctly applying game theory to find their optimal bidding strategy strike me as remote). The smarter bidder could always add a random component back to its bid, to put it in the position it was in without the actuary, so it seems reasonable that it'll be better off. I've done simulations (admittedly with somewhat arbitrary bid distributions) that confirm that there's a slight advantage to being smarter than everybody else.

Another case worth thinking about is when four of the five competitors have an actuary, but the fifth does not. In that case, the simulation suggests that disaster lurks for the non-actuarial bidder: He'll sell more insurance than everybody else, but it will be badly underpriced. Thus, hiring an actuary triggers a kind of arms race. You don't want to be the last to have one.

There are a lot of assumptions in these simulation analyses. To a certain extent, the conclusions depend on the bid distributions, the relative size of the errors versus the mean, and the increase in accuracy resulting from using an actuary. For numerous small homogenous risks, you would expect that there's little additional accuracy an actuary could add at the individual account level. Thus, the reduction in the winner's curse is slight, so it wouldn't be cost-effective to have an actuary.

The Winner's Curse Causes Risk-Averseness

Behavior that mimics risk-averseness is a natural outcome of competition by risk-neutral bidders in auctions. Loosely, we can say that the winner's curse causes risk-averseness. This doesn't mean that the winner's curse is the only source of risk-averseness. Individual underwriters may turn risk-averse in response to the threat of losing their jobs. They may also be reluctant to risk putting the company out of business, even though well-diversified stockholders would be satisfied with the gamble.

I feel that the winner's curse is—or should be—considered the prime source of risk-averseness. Fear of losing your job is an example of what economists term the problems of agency. This occurs when the interests of the individual underwriter don't match up with the stockholder. Ideally, with interests perfectly aligned, these fears would not affect the behavior of the firm.

Here's the argument for how the winner's curse causes behavior that appears risk-averse.

1. The expected winning bid is lower than the mean bid.
2. The difference (bias) gets worse as risk increases.
3. Perfectly rational risk-neutral economic actors will act to offset the bias.
4. Since the bias correction increases as the risk increases, it will appear that risk-neutral actors are risk-averse.

The key assertion is that the bias gets worse as risk increases. Here's an example.

Suppose that two underwriters are bidding on a risk. The true value of the risk is $100, but they don't know this. Their bids are independent and identically distributed and are equal to the mean ($100) plus an error term that looks like this:

Table 1: Distribution of errors in bids

Bid error	–$9	0	+$9
Probability	1/3	1/3	1/3

(Assume that the winner is decided by a coin flip in the event of a tie.) Here, if the underwriters don't adjust their bids to offset the winner's curse, the probabilities of the winning bids are as follows:

Table 2: Distribution of winning bids

Winning Bid	$91	$100	$109
Probability	5/9	3/9	1/9

So the mean winning bid will be ($91 x 5/9) + ($100 x 3/9) + ($109 x 1/9) = $96. If the underwriters want to offset the winner's curse, each must add $4 to his bid.

Now imagine that the errors get twice as big. Table 3 gives the distribution of the errors in the bids.

Table 3: Revised distribution of errors in bids

Bid error	–$18	0	+$18
Probability	1/3	1/3	1/3

It's easy to see (and to compute) that the mean winning bid is now $92. If the underwriters want to offset the winner's curse, each must now add $8 to his bid.

The bidders must make a bigger adjustment to their bids, because their information is worse. The bigger adjustment looks like risk-averseness.

More generally, if you imagine any set of bids as comprising a mean plus random error terms, increasing the variability of the error terms will increase the winner's curse bias.

Properties of Winner's Curse Risk-Averseness

The risk-averseness associated with the winner's curse has some interesting properties. First, it doesn't diversify away. The bias caused by the winner's curse is an expected value phenomenon. Expected values add, no matter how the underlying random variables are correlated.

Second, unlike traditional models of risk-averseness, it's unnecessary to specify a risk measure and it's not necessary to specify a multiplier. For example, one classic risk measure is the standard deviation. The actuary would have to specify a rule such as "add 1.2 times the standard deviation" to convert the risk measure into a dollar charge for risk. The "1.2" comes from a fairly complicated financial model. This kind of model isn't necessary for correcting the bias caused by the winner's curse. The bias correction is solely a function of the joint distribution of bids and the form of the auction.

Next, the risk-averseness doesn't directly depend on the "process risk." In the example above, the correct price of $100 might correspond to a risk process of one $1,000 claim every 10 years. This randomness isn't considered by the winner's curse bias correction. Theoretically, you would get the same winner's curse bias correction if you had a risk process of one $100,000 claim every 1,000 years. Of course, this conclusion depends on the distribution of errors in the bids not being affected by the process risk.

One thing to notice is that the number of bidders makes a difference. An extreme value drawn from a larger population will, on average, be more extreme than one drawn from a smaller population. In the insurance context, the minimum of five bids will, on average, be lower than the minimum of three bids.

The more bidders, the stronger the bias caused by the win-

Table 4: Mean Winning Bid as Function of Standard Deviation

Five bids assumed independently, identically gamma distributed with mean = 100
High bid wins
10,000 iterations

Standard Deviation	Mean Winning Bid	Bias
5	105.9	5.9
10	111.8	1.8
15	118.0	18.0
20	124.2	24.2
25	130.5	30.5
30	136.8	36.8
35	143.2	43.2
40	150.0	50.0
45	156.2	56.2
50	162.9	62.9
55	169.5	69.5
60	176.1	76.1
65	183.1	83.1
70	189.1	89.1
75	196.3	96.3
80	202.7	102.7
85	208.9	108.9
90	216.5	116.5
95	222.8	122.8
100	228.7	128.7

A client with a well-run risk management function almost by definition knows itself better than the other bidders.

ner's curse. But there's another effect, too. Table 4 below shows the results of a simulation. The auction was assumed to have five bidders. Bids were assumed to be independent, identically gamma-distributed, with mean 100. The standard deviation varied from 5 up to 100. Each row gives the results for 10,000 simulated auctions with the indicated standard deviation.

Say that the bidders want to adjust their bids so that the mean winning adjusted bid has the same average as the mean unadjusted bid. For this auction, it's pretty accurate to use the following rule of thumb: "Add 1.2 times the standard deviation of the bid distribution to your bid."

It's interesting to note that for this case, the standard deviation—dear old, much-maligned standard deviation—is almost exactly the right risk measure. Also interesting to note is that the risk-loading factor (the 1.2) is determined by the winner's curse model; it's not another parameter that needs estimation.

William of Occam would like this model.

Anyway, imagine now that there are 400 bidders instead of five. The winning bid will almost certainly come from the extreme tail of the original bid distribution, so a risk measure such as the expected shortfall (of the bid distribution) is more appropriate than the standard deviation. The measure of risk is dependent on the number of bidders.

The auction form makes a difference. Many economics papers discuss different auction forms: the ascending-bid auction (think of Cary Grant in *North by Northwest*); the first-price, sealed-bid auction, where the highest bid wins and the winner pays his bid; the second-price, sealed-bid auction, where the highest bid wins the auction and the winner pays the second highest bid; Dutch auctions, where the price drops according to a set schedule and the first bid wins the auction, etc.

A key result is the "Equivalence Theorem." This says that if all parties act rationally and certain regularity conditions hold, the expected revenue to the seller doesn't depend on the form of the auction. It's true for a variety of auctions.

For personal lines risks and small commercial lines risks, the auction is a winner-takes-all auction. Large commercial risks and large reinsurance treaties are often covered in shares. Different competitors take different shares of the risk.

The price is often on a best-terms basis. Best terms means that if, say, five insurers are bidding on a risk and the client picks the lowest three bidders to share the risk, the rate paid to the successful bidders will be that quoted by the third-lowest bidder, not the lowest bid.

If you assume for the moment that each bidder's bid follows the same continuous distribution, and that that distribution is symmetric around the mean, then in the case described above, there is no bias, because the winning rate is the median of the sample of five bids, and the median is an unbiased estimator of the mean for a symmetric distribution.

One corollary of this for small risks: As the size of a risk shrinks, at some point it becomes uneconomic to have shares, so the auction becomes winner-takes-all. Assume different-size risks are all getting the same number of bids, say, five. The larger risk is paying the third smallest of five bids, the intermediate risk the second smallest of five bids, and the small risk the smallest of five bids. Without some adjustment, the loss ratio for the small risk will be worse than for the larger risks.

For reinsurance and for large sophisticated insureds, there's another factor to consider. The client might decide to retain the risk. In this case, the client is acting like another bidder. A client with a well-run risk management function almost by definition knows itself better than the other bidders. On the other hand, the client won't have access to extensive experience from similar insureds—especially for higher layers—that the other bidders might have. Depending on the layer covered, the client may be better or worse at estimating than the other bidders.

Do Bidders Correct for the Winner's Curse Bias?

It is important to keep in mind that rationality is an assumption in economics, not a demonstrated fact.

—Richard H. Thaler, *The Winner's Curse*

…these paradoxes are of relatively little significance for economics.

—Hirshleifer and Riley, *The Analytics of Uncertainty and Information*

If you're aiming at a target and always hit slightly to the left of it, you can correct your aim. In the same way, if an estimator is biased, you can undo the bias by adding a correction to the estimate. Some economists argue that this is what happens in real-world auctions: The bidders adjust their bids so that the expected winning bid has no bias, or more generally, that the bid is optimal in some game-theoretic sense.

Other economists argue that experimental evidence shows people don't make bias corrections. It's also interesting to note the discrepancy between theory and practice regarding the effect of increasing the number of bidders. Adding bidders to the auction increases the winner's curse bias. Undoing the bias means adding more back to the original bid. This theory is clear, but it contradicts what people do in practice. In the face of increased competition, they shave their margins.

One argument given against the relevance of the experimental results is that the experiments are usually conducted on college students, who lack real-world experience. People who have a long-term interest in a particular kind of auction should do better.

On the other hand, it's hard to imagine that people correctly infer the correct game-theoretical bidding strategy in most bidding situations. After he measures your living room, the contractor goes back to his office and integrates the joint bid distribution to find his optimal bidding strategy.

The argument about increasing the number of bidders may fail because the model is too simple. Modifying the model to reflect fixed expenses might give a different result, one more in line with our experience and common sense.

My feeling is that experienced bidders—even those with no concept of probability and statistics—will eventually notice that there is something wrong and make an adjustment. The adjustment might not be exactly correct, but it will be in the right direction.

The Importance of Being Credible

All you need in this life is ignorance and confidence; then success is sure.

—Mark Twain (1835–1910), Letter to Mrs. Foote, Dec. 2, 1887

Imagine that you are the actuary at a personal lines company and you're creating prices by territory for personal auto. Your company's volume varies by territory. In some territories, you have a large volume of data; in others, you have scanty data. What happens if you don't apply actuarial credibility techniques to your data?

Well, in the real world, you'd be fired. But assume not, for the time being. Also, assume that getting the loss costs right is all there is to ratemaking—everybody loads the same profits and expenses onto the loss costs—and you don't compare your proposed rates to competitors' before finalizing them. What will happen, given these suppositions?

In the territories where your company has a lot of data, you'll come up with an estimate of the territory loss costs that's close to the correct average. In the territories with little data, your estimates will randomly be either much too high or much too low. Where your estimates are too high, you'll sell even less insurance than before. Where your estimates are too low, you'll sell lots of insurance. And the next time you do your territorial review, you'll wish you hadn't.

Whether you sell much insurance in the territories where your rates are pretty accurate depends on the state of the market. If there are a lot of ignorant competitors, then knowing the right answer will avail you little—you won't be able to sell insurance in those territories. This suggests a scenario where your company grows in those territories where it's ignorant, and shrinks in those territories where it knows something. A weird oscillation ensues.

What happens in reality is that actuarial credibility techniques are applied. This dampens the swings in the rates. Also, you can compare your rates against your competitors' and adjust your rates in response to what you see. This can tame the winner's curse.

Conclusion

The winner's curse helps us understand many things about insurance pricing. But there are still some areas where the theory and practice are not in line. There's a decision to make. The winner's curse seems to demand a form of risk load. Is this risk load in addition to the risk load commonly used by actuaries now? Or is it a (partial or complete) replacement for the risk load we use now?

The theory is incomplete, and the inconsistencies between this and traditional risk load methods haven't been worked out. In the meantime—what do we do? ●

CHRISTIAN SVENDSGAARD is a senior actuary at Swiss Re in Armonk, N.Y.

Further Reading

Thaler, Richard H., The *Winner's Curse: Paradoxes and Anomalies of Economic Life*, Princeton University Press, 1994.

Capen, E. C., Clapp, R. V., and Campbell, W. M.: "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology*, June 1971.

Cox, James C., and Isaac, R. Mark: "In Search of the Winner's Curse," *Economic Inquiry*, October 1984.

EMPLOYMENT CPS SERVICES NETWORK

YOU WANT THE BEST FOR YOURSELF ...

YOU DESERVE HELP IN GETTING IT.

THE ACTUARIAL MARKETPLACE IS OUR SPECIALTY

We are one of the largest recruiting firms in the country. Our reputation for service is unsurpassed. We have been in the recruiting business for almost 30 years—let our experience work for you.

CALL ONE OF OUR RECRUITING SPECIALISTS TODAY

PATRICK KILCOYNE
KRISTINE FORNEK
One Westbrook Center #600
Westchester, IL 60154-5799
(708) 531-8289
Fax (708) 531-8368
pkilcoyne@cps4jobs.com

MARY O'CONNELL
ALAN IRISH
50 Federal Street #301
Boston, MA 02110-2585
(617) 368-3550
Fax (617) 368-3562
mary@cpsboston-jobs.com