

Pricing & revenue architecture

What is the winner's curse?

The winner's curse is the shortfall that can follow winning a bid on a common, uncertain value: the winning bid tends to rest on the largest overestimate.

Author	Published	Reading time	Words
Sinan Isoglu, MBA (Quantic) Commercial growth leader, lecturer and doctoral researcher	18 October 2026	18 min	4,014

Management summary The winner's curse is the shortfall between what the winner of a competitive bid expected and what the item turns out to be worth, when every bidder estimates the same uncertain value and the winning bid tends to rest on the largest overestimate. This concept guide defines it (Capen, Clapp and Campbell; Bazerman and Samuelson; Thaler), explains why winning is itself information about the estimate, reports the oil lease sales and a classroom experiment in which the average winning bid for items worth US\$8 was US\$10.01, sets out why bidders keep falling for it, and gives a bid record for a firm that prices contracts on its own cost estimate, labelled as the author's and untested. The sources are an industry modelling paper, one classroom experiment and one review column.

Contents	§ 1	What is the winner's curse?	2
	§ 2	Why is winning itself a clue about the estimate?	3
	Figure 1	More bidders, higher winning estimates, a lower bid needed	5
	§ 3	What did the oil lease sales show?	6
	§ 4	What happened when MBA students bid for jars of coins and paper clips?	7
	§ 5	Why do bidders keep falling for it?	8
	§ 6	What can a firm that bids for contracts carry over?	9
	Figure 2	A bid record	10
	§ 7	How far do oil leases and jars of coins travel?	11
		References	13

What is the winner's curse?

§ 1 of 7

It is the bad news that can come with winning. Bazerman and Samuelson (1983)¹ define it in their first sentence: the winner's curse "occurs in competitive situations when a successful buyer finds that he or she has paid too much for" an item of uncertain value. Researchers before them had suggested that the winner of a sealed-bid auction "will often lose", that is, "the object acquired will be worth less than the price paid".

¹ Bazerman, M. H., & Samuelson, W. F. (1983). I won the auction but don't want the prize. *Journal of Conflict Resolution*, 27(4), 618-634.

Thaler (1988)³ sets out the case in which the idea began. Several oil companies bid for the drilling rights to one parcel: "Let's assume that the rights are worth the same amount to all bidders, that is, the auction is what is called a common value auction." Each firm has its experts estimate that value, and the estimates are unbiased: "the mean of the estimates is equal to the common value of the tract." Because oil is hard to estimate, the estimates scatter, the firms with high estimates tend to bid more, and the winner may be the firm whose experts guessed highest. "If this happens, the winner of the auction is likely to be a loser."

³ Thaler, R. H. (1988). Anomalies: The winner's curse. *Journal of Economic Perspectives*, 2(1), 191-202.

Thaler separates two versions. In the first, "the winning bid exceeds the value of the tract, so the firm loses money"; in the second, "the value of the tract is less than the expert's estimate so the winning firm is disappointed", which "can apply even if the winning bidder makes a profit, as long as the profit is less than expected at the time the bid was made." He traces the concept to "three Atlantic Richfield engineers", Capen et al. (1971)², whose own summary is that "a lease winner tends to be the bidder who most overestimates reserves potential". They wrote about sealed-bid lease sales and held that the phenomenon "pervades all competitive situations".

² Capen, E. C., Clapp, R. V., & Campbell, W. M. (1971). Competitive bidding in high-risk situations. *Journal of Petroleum Technology*, 23(6), 641-653.

In this guide, the winner's curse is the shortfall between what the winner of a competitive bid expected and what the item turns out to be worth, which arises when bidders bid on uncertain estimates of a value that is the same for all of them, so that the winning bid tends to rest on

the largest overestimate. It is a tendency, not a rule for every win: Capen et al.² say they are describing trends, not what will happen each time a tract is bought. It is also a mistake by bidders, not a property of auctions as such: Thaler, citing Cox and Isaac (1984), notes that it “cannot occur if all the bidders are rational”. And it concerns objective value. Bazerman and Samuelson¹ point out that a bidder can knowingly pay more than an item’s objective value without regret when it has “personal or intrinsic value (e.g., a painting)”, and the central theme of their study is “the existence of an objective winner’s curse”. On this guide’s reading, the tracts a bidder wins are a sample selected by the bidder’s own estimation errors, which makes the curse a case of what the [selection bias guide](#) describes: a path into the sample that depends on something that matters for the outcome.

Why is winning itself a clue about the estimate? § 2 of 7

Because the bid that wins is usually the one built on the highest estimate. Bazerman and Samuelson¹ give the reasoning in three steps: while the average bidder may estimate the value accurately, some bidders overestimate it and some underestimate it; “the bidder who most greatly overestimates the value of the commodity will typically win the auction”; and “the amount of overestimation will often be greater than the difference between the winning bidder’s estimate of the value of the commodity and his or her bid”. Even if bidders bid below their estimates on average, “a winning bid drawn from the right tail of the bid distribution may exceed the actual value of the good”.

Capen et al.² put the same point in terms of a firm’s record. A bidder’s estimates can be right on average across all the tracts it looks at and still be wrong on average across the tracts it wins, because it has a poor chance of winning where it underestimated the value and a good chance where it overestimated it. So the bidder “tends to win a biased set of tracts”, namely “those on which he has overestimated value or reserves”. Only where there is no competition do the high and the low estimates offset each other.


The correction is an inference that bidders tend to skip. In Bazerman and Samuelson’s¹ words: “If an individual assumes that his or her bid will win the auction, this piece of data should indicate that the bidder has probably overestimated the value of the commodity in comparison to other competitors.” Thaler states what a rational bid requires: “Rational bidding requires first distinguishing between the expected value of the object for sale, conditioned only on the prior information available, and the expected value conditioned on winning the auction.”

More competitors pull in two directions, at least in a first price auction, where the highest bidder pays its bid. “An increase in the number of other bidders implies that to win the auction you must bid more aggressively, but their presence also increases the chance that if you win, you will have overestimated the value of the object for sale”, Thaler writes. Capen et al.² drew three rules from simulating lease sales, rules they say are not without exceptions but seem to apply at the normal level of competition and with the large uncertainties behind their value estimates: the less information a bidder has compared with its opponents, the more uncertain its own estimate and the more bidders there are above three, the lower it should bid. In the analytical example in their appendix, where value estimates follow an exponential distribution and every opponent bids half its estimate, the bid level that maximises the expected present worth of a bid rises to a maximum of 0.28 of the bidder’s own estimate at two opponents and then falls, to 0.15 at 12 opponents, because the tracts a bidder wins tend to be those on which it overestimated the value.

Capen et al.² also worked through a small hypothetical sale. Ten tracts, all assumed productive, are each worth US\$10 million at a 10 per cent discount rate, every bidder estimates the total of US\$100 million correctly but misjudges single tracts, and every bidder bids half its estimate. With one bidder the return is about 17 per cent after tax. As rivals join, each tract goes to the bidder with the highest estimate of it, and the industry’s return falls to 14 per cent with two bidders, about 11 per cent with four, about 8 per cent with eight and about 6 per cent with sixteen. Figure 1 shows the table that closes the example.

Figure 1 More bidders, higher winning estimates, a lower bid needed

The ten tracts are worth US\$100 million whatever the number of bidders. With sixteen bidders the winning estimates add up to US\$325 million, and each bidder would have had to bid 0.31 of its estimate for the industry to earn about 10 per cent.



Every bidder in the example estimates the total correctly and bids half its estimate; the industry's return is about 17 per cent with one bidder and about 6 per cent with sixteen. The bid levels assume that all bidders use the same return criterion and the same strategy. The right-hand bars are not the bids made, which stayed at half of each estimate; they are the bid level at which the industry would have earned about 10 per cent.

Source Capen et al. (1971), p. 645, the unnumbered table that closes their hypothetical ten-tract sale: for 1, 2, 4, 8 and 16 bidders, the total of the highest value estimates on the ten tracts and the bid level, as a fraction of a bidder's own value estimate, that would have given the industry a 10 per cent return. Every tract is worth US\$10 million, the ten together US\$100 million. The authors give the bid levels as what their example implies and say they are valid for that example only.

The value of the tracts never changes; what changes is which estimate wins. The authors note that their bid levels hold only for this example, in which every bidder uses the same return criterion and the same bidding strategy, and that companies in the real world “are not so inclined to play that way”. The direction, they argue, is general: falling returns as bidders are added “appears to us a general rule of sealed bidding”.

What did the oil lease sales show?

§ 3 of 7

Capen et al.² start from the industry's record in sealed-bid lease sales. In the Gulf of Mexico, most analysts had found that "the industry probably is not" earning the return on its investment that it intended, and, leaving out the era before 1950, "the Gulf has paid off at something less than the local credit union." "It is our view", the authors write, that none of the usual explanations (poor estimates of the oil and gas in place, overly conservative cost estimates, misjudged production allowables, timing or bad luck) was the major cause, though all may have contributed. Their explanation is the winner's curse. On this guide's reading, that is an inference: the paper cites the return studies rather than reporting them, and it offers no test that separates the curse from the other causes.

What the paper does show is how far apart serious bidders' bids were, a spread the authors read as divergent value estimates; bids also carry differences in discount rates and bid levels, they note, and Thaler adds that some low bids may be placed in the hope that no one else will bid. In the paper's Table 1, the highest bid was 10, 7, 109 and 26 times the lowest serious bid on one tract in each of four sales between 1967 and 1969; the ratio between the highest and the lowest serious bid, the authors write, runs to as much as 100 and is commonly 5 or 10. In the 1969 Alaska North Slope sale, Atlantic Richfield and Humble, partners in much exploration and so probably holding much the same information, bid independently on 55 tracts. Humble's bid ran from about 0.03 of Atlantic Richfield's to about 17 times it, and no one had been able to identify "any pattern or hint of correlation in these numbers".

Figures on the North Slope bids, which the authors credit to an oral communication from Lawrence Friedman at a 1969 seminar, also show how far the winners stood above the next bidder. The winning bids added up to US\$900 million and the second-highest bids to US\$370 million, so the runner-up was willing to bid, on average, 41 per cent as much as the winner. On 26 per cent of the tracts the second-highest bidder would not have won even with four times its bid, and if every winner had bid only two-thirds as much, the winners would still have kept 79 per cent of the tracts they won. From this the authors conclude that lowering bids to reach a target return does not drastically reduce a firm's chances of winning acreage.

Later work, as Thaler (1988)³ reports it, is mixed. In his account, Mead et al. (1983) calculated returns on 1,223 Gulf of Mexico leases issued between 1954 and 1969 and found an average present-value loss of

US\$192,128 per lease at a 12.5 per cent discount rate; 62 per cent of the leases were dry and only 22 per cent profitable, and Thaler reads this as consistent “at least with version 2 of the winner’s curse”. Thaler also reports that Hendricks et al. (1987), with a 5 per cent real discount rate and other assumptions, found that firms would have made profits even at constant real oil prices, yet for 12 of the 18 firms that bid often, the multiple of their own bids that would have maximised profit, with all other bids unchanged, was below one, with a median of 0.68, which Thaler reads as “some support for the winner’s curse”. Bazerman and Samuelson¹, writing in 1983, called the field evidence “sketchy at best” and noted 1970s estimates that the US government had received approximately the full monetary value for the offshore tracts it sold.

What happened when MBA students bid for jars of coins and paper clips? § 4 of 7

The winning bids exceeded the value on average. Bazerman and Samuelson (1983)¹ ran the experiment with 419 MBA students in 12 microeconomics classes at Boston University. Each class held four sealed-bid auctions, each for a different commodity, and, unknown to the students, every commodity was worth US\$8.00: 800 pennies, 160 nickels, 200 large paper clips valued at four cents each and 400 small paper clips valued at two cents each. The students “were told that the highest bidder would pay his or her bid” and receive the stated value, so “subjects bid on the value of the commodity, not the commodity itself”. The students also gave a best estimate with 90 per cent confidence bounds, a US\$2.00 prize went to the closest estimate in each auction, and no feedback was given until the end. Each student was told the size of the auction group for each auction, between 4 and 26 bidders, but no student was assigned to a group: for each auction the authors averaged the highest bid over all possible groups of that size in the class, and 12 classes with four auctions each gave 48 auctions. The paper therefore reports the payment rule the students were told, not any payment made.

The students underestimated the value: their mean estimate was US\$5.13, US\$2.87 below it, which should have worked against the curse. Even so, “The mean AWB for the 48 auctions was \$10.01, with a standard deviation of \$5.48”, where AWB is the average winning bid, so the average auction implied a loss of US\$2.01 for its winner. The average winning bid was between US\$8.00 and US\$9.00 in 3 auctions and over US\$9.00 in 23, so it exceeded the value in 26 of the 48.

The two factors the authors studied went with the winning bid as they predicted, though weakly. Neither the uncertainty about a commodity,

measured by the width of the students' 90 per cent ranges, nor the number of bidders significantly affected the mean bid in an auction; across individual bids, the effect of uncertainty was significant but small, a dollar more of range going with a bid lower "by only three cents". The average winning bid did rise with uncertainty and with the number of bidders, but the regression explained "only about 18% of the variation", and the authors ask that its coefficients be read as "no more than suggestive". On those coefficients, starting from the means (a 90 per cent range of US\$5.80 and 15 bidders), a 20 per cent rise in uncertainty goes with a 43 per cent higher average winning bid and a 20 per cent rise in the number of bidders with a 5 per cent higher one, and the most uncertain commodity was predicted to produce a loss for the winner with as few as 4 bidders.

Why do bidders keep falling for it?

§ 5 of 7

In the classroom, because bids followed estimates and little else. Bazerman and Samuelson¹ write that the evidence suggests "subjects employ naive bidding strategies, basing bids upon unconditional value estimates and disregarding relevant information such as the uncertainty surrounding the commodity and number of bidders", and that they did not adjust their bids "with changes in the number of competitors". They had predicted this for two reasons: bidders would miss the inference to be drawn from having the highest bid, and, as the authors put it, bidders facing many rivals commonly reason "I will have to bid closer to the real value (my estimate)" if they are to win.

Capen et al.² describe the same reasoning in the oil industry. Most people "assume that the tougher the competition" the more they must bid to stay in the game; that is right, the authors reply, if the aim is to buy the most acreage, and probably wrong if the aim is a given rate of return. Nor does being right on average protect a bidder, because it tends to win when it most overestimates value. Their conclusion is that "Unexpectedly low rates of return, however, follow the industry into competitive lease sale environments year after year", although each competitor thinks it is playing a reasonable strategy.

Laboratory work reviewed by Thaler (1988)³ points the same way. As Thaler reports them, the common value auctions run by Kagel and Levin (1986) typically gave small groups of three to five bidders profits that averaged 65.1 per cent of those predicted for risk-neutral bidders in equilibrium, while groups of six or seven lost US\$0.88 per auction period where that prediction was a profit of US\$4.68; "the subjects bid more aggressively in the larger groups", where the prediction requires bidding more conservatively. He also reports that in laboratory

low-price auctions run by Dyer et al. (1987), in which both small and larger groups lost money, construction firm managers, who bid for contracts in their work, “did no better or worse than the students”, and that the experimenters believe the managers had learned rules of thumb that avoid the curse in their own markets but could not be applied in the laboratory.

Avoiding the curse also has a price. Thaler notes that a firm that understood it would face a choice: “If you react by optimally reducing your bids, then you will avoid paying too much for leases, but you will also win very few auctions.” Capen et al.² argue that the loss of volume is smaller than it looks: the estimates diverge so widely that a bidder’s chance of winning depends more on its reserves estimate than on its bid level, as the North Slope figures above suggest.

What can a firm that bids for contracts carry over?

§ 6 of 7

The sources do not study contract bids between firms, so this section is the author’s reasoning. Thaler describes the mirror case in a footnote: “In a low price auction, such as construction contract bidding, the winner is the one who submits the lowest bid.” Bazerman and Samuelson¹ note that in competitive procurements “it is commonly held that the contract winner may be simply the most optimistic firm (the one that most grossly underestimates program costs) rather than the most efficient”, and they add that frequent cost overruns stem partly from such misestimates “but are frequently due to poor contract incentives for the winning firm to seek cost economies”.

On this guide’s reading, a seller that bids a fixed fee on its own cost estimate meets the curse in mirror image. Where several bidders face the same uncertain scope of work, the lowest bid tends to rest on the lowest estimate of that scope; where their costs genuinely differ, a low bid can simply come from a cheaper bidder, which is the distinction, in the common view Bazerman and Samuelson¹ report, between the most optimistic firm and the most efficient one. The same reading applies to a buyer choosing the lowest of many bids, with their caution that an overrun can come from contract incentives as well as from the estimate. Reverse auctions, a buying format in which suppliers bid against each other, are not covered here.

Bazerman and Samuelson¹ end with “a simple starting point for decision makers”: “Determine if new information would be available if you assumed you were going to ‘win’ the competition.” They offer it as a first answer to a question they put for negotiation and competitive

bidding alike, “how can an individual avoid a transaction in which he or she pays more than the agreement is worth?”, and test it in neither. Figure 2 turns their question, and two of Capen et al.²’s three rules, on the uncertainty of one’s estimate and on the number of bidders, into a record a bidder can keep for each bid: the estimate and its range, the number of serious rivals, what winning would say, the bid against the estimate and, once the work is done, the outcome against the estimate. Counting serious rivals is the kind of inference from thin signals that the competitive intelligence essay discusses.

Figure 2 A bid record

Write down the estimate and its range, the rivals, what a win would say and the bid before the bid goes in, and the outcome after the work is done.

Estimate	Rivals	If we win	Bid	Outcome
Value or cost, with low and high?	How many serious bidders?	What would winning say about the estimate?	Bid against the estimate, and why?	Actual figure against the estimate?

Read the outcome column for the bids you won; that is the set the curse selects.

One row per bid. Fill in the first four fields before the bid goes in and the last after the work is done; compare outcomes with estimates on the bids that won.

Source Author’s decision worksheet. The estimate and rivals fields follow the second and third of the three rules in Capen et al. (1971), p. 644, and the first, on information relative to rivals, has no field; the question in the third field is Bazerman and Samuelson (1983), p. 634; the bid field follows Capen et al. (1971), pp. 644 to 645 and 653; the outcome field follows Capen et al. (1971), p. 643, on estimates that are right on average but not on the tracts won. The fields are reader inputs; no company data is supplied.

An illustration, with no company data: a services firm bids a fixed fee for a data migration against four other firms that have read the same tender documents. Its estimate is 1,200 hours, with a range of 800 to 2,000 hours, because the state of the client’s data is unknown to every

bidder. Before bidding, the record asks what a win would say: with four rivals reading the same documents, a winning fee is more likely to rest on an hours estimate at the low end of the bidders' range than on an efficiency the firm can name. The firm can then price from a higher point in its own range, or bid with a stated assumption about the data that reopens the scope if it proves wrong. After each project it wins, the record sets the hours used against the estimate. Capen et al.²'s distinction is the reason to look at the won bids: estimates that are right on average across all bids can still be too low on average across the bids that won. Where the firm also does work without competition, the same comparison there shows whether its estimates run low everywhere or only on the bids it wins, the contrast Capen et al. draw between competitive and noncompetitive settings. A won bid is an outcome, not yet a result; the win-loss essay makes the neighbouring point that an outcome is not yet its reason, and keeps the two apart. These are the author's examples and have not been tested.

How far do oil leases and jars of coins travel?

§ 7 of 7

Not far without care. Capen et al. (1971)² is a paper by three engineers at one oil company, written, as they say, to share their thoughts on the theory of competitive bidding; its evidence is the spread of published bids, second-bid figures for one sale taken from an oral communication, a hypothetical ten-tract example, a simulation model, an analytical appendix and return studies it cites but does not report. Bazerman and Samuelson (1983)¹ ran one classroom experiment with MBA students, as an introduction to decision making under uncertainty, without feedback, with auction groups computed after the fact and with a payment rule that was announced but is not reported as carried out; they treat the students' underestimation as specific to their task and leave "The form of training that can lead to optimal improvement in bidding judgments" to further work. Thaler (1988)³ is a review column: its field evidence on later lease returns comes from other authors and points both ways. None of the three studies contract bidding between firms.

On this guide's reading, two things carry over with some confidence. When bidders estimate a common value with error, the bid that wins tends to rest on the largest overestimate, as long as bids rise with estimates: that follows from the selection logic all three sources share, the students in the one experiment this guide draws on did not correct for it, and the field returns are consistent with it though mixed. And the record of won bids is the place to look for it. How large the effect is for a firm bidding on contracts, and whether experience removes it, the

sources do not say. The construction managers in the laboratory study Thaler reports did no better than students, yet construction firms bid for contracts all the time and, as he notes, “would soon go bankrupt if they fell prey to the winner’s curse”.

 End of essay

References

- 1** Bazerman, M. H., & Samuelson, W. F. (1983). I won the auction but don't want the prize. *Journal of Conflict Resolution*, 27(4), 618–634. <https://doi.org/10.1177/0022002783027004003>
- 2** Capen, E. C., Clapp, R. V., & Campbell, W. M. (1971). Competitive bidding in high-risk situations. *Journal of Petroleum Technology*, 23(6), 641–653. <https://doi.org/10.2118/2993-PA>
- 3** Thaler, R. H. (1988). Anomalies: The winner's curse. *Journal of Economic Perspectives*, 2(1), 191–202. <https://doi.org/10.1257/jep.2.1.191>

How to cite this essay

Isoglu, S. (2026, October 18). What is the winner's curse? Sinan Isoglu. <https://isoglu.com/insights/journal/what-is-the-winners-curse/>
Online version: <https://isoglu.com/insights/journal/what-is-the-winners-curse/>

Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher

© 2026 Sinan Isoglu · isoglu.com · All rights reserved