Prediction errors and the winner's curse.*

Olivier Compte†

September 2004

Abstract

A popular explanation for the winner's curse is that in correlated or interdependent value environments, bidders fail to take into account the information on other's estimate (hence, given the correlation or the interdependence, on own valuation) conveyed by winning.

Another explanation, initially proposed by Capen and al. (1971), is that bidders make estimation errors, and that competition induces a selection bias in favor of most optimistic bidders.

The main purpose of this paper is to show that these explanations are not equivalent. In particular, the latter one extends to settings in which values and estimates are drawn from independent distributions, while the first does not.

The paper also discusses the role of *over-confidence* in the accuracy of own signals in explaining the winner's curse and its persistence: we find that both explanations build on errors that may have their roots in that same cognitive bias.

Key words: Winner's curse, private values, prediction errors, overconfidence.

---

*This paper builds on Section 2 of "The winner's curse with independent private values". I thank Philippe Jehiel, Eric Maskin, Ran Spriegler, Jean Tirole, Shmuel Zamir as well as seminar participants at CORE, ESSET 2001 (Gerzensee), the Institute for Advanced Studies (Princeton), the University of Pennsylvania, Université d'Aix-Marseille (Greqam) for helpful comments.

†CERAS-ENPC, e-mail compte@enpc.fr.

1. Introduction

In some competitive environments, "successful" bidders are not so successful after all: "successful" bidders (that is, those who won the competition) tend to obtain returns that (on average) lie below initial projections. This discrepancy between realized and anticipated returns, and the possibility that winning bidders end up making losses, has been called the *winner's curse*.[1]

The most popular explanation for the winner's curse is based on bidders' failure to incorporate the information conveyed by winning into their bidding strategy. Winning typically conveys information on *others'* value estimates. Thus, when valuations are interdependent or drawn from correlated distributions, winning conveys information on one's *own* valuation. If bidders fail to take into account this information, they end up with a biased estimation of their valuation.

Another explanation, initially proposed by Capen and al. (1971) in the context of competition for oil fields, is that bidders make estimation errors, and that competition induces a selection bias in favor of most optimistic bidders. Capen and al. (1971) consider the problem of bidding for a tract that has the same (unknown) value for each bidder, and they argue that:

> "in competitive bidding, the winner tends to be the player who most overestimates true tract value.... [So] a player tends to win a biased set of tracts - namely those on which he has over estimated value or reserves".

Do these explanations differ, or are they just two different ways to say the same thing?

Looking at the many textbooks and academic papers on the topic, one could be tempted to opt for the latter view. Indeed, what is most often retained from Capen and al.'s example is that bidders face a *common value* environment, and that they fail

[1] For an introduction to winner's curse, see for example Milgrom (1989) or Thaler (1992, chapter 5). For evidence of the winner's curse in experiments, see Bazerman and Samuelson (1983) and Kagel and Levin (1986). See also Kagel and Levin (2002) for a recent survey of experimental evidence.

to incorporate the information conveyed by winning into their bidding strategy; hence the curse.

This paper will argue that these explanations do differ. In a nutshell, we will show that while the first explanation relies on the fact that bidder's valuation are correlated -or interdependent, Capen and al's explanation does not.

Intuitively, the reason is that competition tends to select optimistic bidders, whether bidders' valuations are drawn from independent distributions or not. One could for example re-write Capen and al.'s argument as follows:

> In competitive bidding, a bidder is more likely to win when he over-estimates his own value for the object. So a bidder tends to win a biased set of objects - namely those on which he over-estimates his own value.

The next Section proposes a model that will formalize that intuition. The main feature of our model is that it allows for prediction or estimation errors, which may for example result from players being *overconfident in the accuracy of their own signals.* Because of these predictions errors, bidders may either make optimistic or pessimistic predictions. The point is that even if on average these errors cancel out, competition induces a selection bias in favor of optimistic bidders. We show in Section 3 that this selection bias exists even when valuations are drawn from independent distributions, hence even when winning conveys no information on one's own valuation; and failing to take into account this selection bias is potentially harmful to bidders.

Finally, in Section 4, we analyze a common value setting and characterize the two possible sources of curse: the standard one, based on the agent's failure to take into account the information on other's estimates conveyed by winning, and the second one, based on the fact that competition tends to select optimistic bidders. These two effects are compared in a simple example. We then discuss more generally the role of *overconfidence in own signals* in explaining the winner's curse and its persistence.

2. A model with prediction errors.

Our first objective is to present Capen and al.'s insight in the simplest way, and show that this insight is valid in a private value setting. Our model builds on an example given in Milgrom (1989).[2]

Consider the problem of bidding for construction contracts. There are n potential contractors bidding for a construction contract. Each contractor $i = 1, ..., n$ has a cost C_i of doing the job. Contractors however do not know precisely what the job will cost, but they get an estimate X_i of C_i.

The process by which contractor i gets the estimate X_i is not modelled. To an outside observer however, the costs C_i and the estimates X_i can be regarded as random variables. For simplicity, we assume that

$$X_i = C_i + \widetilde{\varepsilon}_i,$$

where $\widetilde{\varepsilon}_i$ is an estimation error. The errors $\widetilde{\varepsilon}_i$, $i = 1, ..., n$ are assumed to be independent across bidders, independent of costs and satisfy:

$$E[\widetilde{\varepsilon}_i] = 0.$$

So the estimate X_i is unbiased on average.

In this Section and the next one, we focus on the *private cost* case, in which the individual costs C_i are drawn from independent distributions. We also assume that the distributions over costs and errors are non-degenerate. Formally, we assume that the cost C_i and the error $\widetilde{\varepsilon}_i$ each admit a density (denoted respectively f_i and g_i), that the support of each density is an interval, and that each density is positive on its support.

The *symmetric case* will refer to situations where the marginal distributions over C_i are identical across bidders, and where the errors are drawn from identical distributions.[3]

We now turn to the main assumption of our model.

[2] Milgom (1989) considers a common value example, that we transform into a private value example.
[3] That is, $f_i = f$ and $g_i = g$.

In a standard Bayesian model, contractors would be assumed to know the joint distributions over the costs C_i and estimates X_i. In such a model, each contractor i would be able to compute, for each realization x of X_i, the conditional distribution over costs given x, hence also the conditional expectation of costs

$$Y_i^*(x) \equiv E[C_i \mid X_i = x].$$

Here, in contrast, we shall assume that each agent is aware that costs are distributed independently, but that, based on X_i, each bidder i forms a *possibly erroneous* prediction $\widehat{Y}_i$ of the cost C_i.[4,5] One particularly simple case to analyze is one where the agent takes his estimate at face value, without realizing that he makes estimation errors:

Assumption 1. $\widehat{Y}_i \equiv X_i$.

One possible interpretation of Assumption 1, which echoes numerous work in psychology on overconfidence, is that contractor i believes that the estimate X_i has more predictive content than it really has.

Comment: Of course, other specifications are plausible. For example, less extreme forms of overconfidence would consist in assuming that contractor i (erroneously) believes that

$$X_i = C_i + \lambda \varepsilon_i \tag{2.1}$$

where $\lambda \in [0, 1)$.[6] The interpretation of λ is that, although contractor i realizes that his estimate is subject to errors, he *downplays the magnitude of his own errors* (by a factor λ). In that case, the prediction of costs would be:

$$\widehat{Y}_i \equiv E^\lambda[C_i \mid X_i],$$

[4] The prediction $\widehat{Y}_i$ may thus differ from the Bayesian prediction Y_i^*.

[5] Alternatively, we could assume that each bidder forms a possibly erroneous belief (i.e. conditional distribution) over his cost C_i given his observation X_i. The prediction $\widehat{Y}_i$ should then be thought of the expected value of cost conditional on X_i, taken under that possibly erroneous conditional distribution.

[6] This type of overconfidence is analogous to the one that appears in the finance literature, in which traders are assumed to be overconfident in the informative content or accuracy of their signals. See for example De Long et al. (1991), Kyle and Wang (1997), Odean (1998), or Daniel, Hirshleifer and Subrahmanyam (1998). These papers also discuss the relevent literature in psychology.

where the superscript λ stands for the fact that the expectation is taken assuming that the joint distributions over X_i and C_i is characterized by (2.1).[7]

Another plausible assumption would be that bidders make *computation errors* in computing conditional expectations. One possible specification would then be:

$$\widehat{Y}_i \equiv Y_i^* + \eta_i,$$

where η_i is an independent random variable with zero mean.[8]

Because contractor i makes prediction errors, he will sometimes be too optimistic about his cost, meaning that

$$\widehat{Y}_i < Y_i^*(X_i),$$

and he will sometimes be too pessimistic about his cost, meaning that

$$\widehat{Y}_i > Y_i^*(X_i).$$

The difference between the correct and the actual prediction is denoted H_i:

$$H_i(X_i, \widehat{Y}_i) \equiv Y_i^*(X_i) - \widehat{Y}_i,$$

and, for any given realization of X_i and $\widehat{Y}_i$, this difference will interpreted as the agent's *optimism.* Note that under Assumption 1, bidder i is not optimistic on average, because by assumption $EX_i = EC_i$, hence prediction errors cancel out:

$$EH_i = 0.$$

[7] To fix ideas, assume that C_i and ε_i are drawn from normal distributions, respectively $N(c_0, \sigma^2)$ and $N(0, \nu^2)$. Then

$$\widehat{Y}_i(X_i) = \mu(\lambda)X_i + (1 - \mu(\lambda))c_0$$

with $\mu(\lambda) = \sigma^2/(\sigma^2 + \lambda^2\nu^2)$. So the smaller λ, the more weight contractor i puts on X_i relative to the prior information c_0.

[8] Bidders may fail to correctly compute conditional expectations even when they are given the structure of the model and know all prior distributions. This has been checked experimentally in recent and independent work by Goere and Offerman (2002). Their interpretation for these errors is that bidders fall prey to the base rate fallacy, and they notice the possibility of a curse induced by these errors, which they call a new's curse.

Finally, we assume that contractors participate in a second price auction.[9] Each contractor i simultaneously submits a bid b_i: the winner is the bidder with the lowest bid, and he gets paid a price equal to the second lowest bid for doing the job.

In the private value setting we analyze, it is a dominant strategy to bid one's own cost prediction $\widehat{Y}_i$, so we have:

$$b_i(\widehat{Y}_i) = \widehat{Y}_i$$

To complete the description of the model, we define formally what we mean by winner's curse. We let Δ_i denote the expected difference between realized and predicted costs, conditional on winning, that is,

$$\Delta_i = E[C_i - \widehat{Y}_i \mid i \text{ wins}]$$

Definition 1. *We shall say that contractor i is subject to the winner's curse when* $\Delta_i > 0$.

In other words, the winner's curse refers to situations in which winners underestimate costs: on average, winners are disappointed when the actual cost is realized. Note that Δ_i can be rewritten as

$$\Delta_i = E[Y_i^* - \widehat{Y}_i \mid i \text{ wins}] = E[H_i \mid i \text{ wins}],$$

hence Δ_i can also be interpreted as bidder i's *expected optimism*, conditional on winning. It thus follows immediately that in a standard Bayesian model there can be no winner's curse, since then, $H_i(X_i, \widehat{Y}_i) = 0$ with probability one.

[9] We focus on the second price auction, so that we need not worry about the bidder's beliefs about other bidders' bids. However, our analysis can be easily extended to first price auctions, assuming for example that in equilibrium, bidders know (or have learned) the distribution over their opponent's bids. Under Assumption 1, $b_i(\widehat{Y}_i)$ would solve,

$$b_i(\widehat{Y}_i) = \arg\max(b_i - \widehat{Y}_i)\Pr(b_i < \underline{b}_{-i}).$$

3. Main result.

In this Section, we show that under Assumption 1, bidder i is subject to a winner's curse (even though costs are drawn from independent distributions) and that bidders may end up making losses when competition is fierce enough. We will also examine how increased competition affects the curse.

3.1. The winner's curse.

If bidder i were certain to obtain the contract, then on average, his prediction errors would cancel out and he would not be subject to the winner's curse. Formally, we would have:

$$\Delta_i = E[H_i \mid i \text{ wins}] = EH_i = 0.$$

When he competes with other bidders, bidder i is (typically) no longer certain to obtain the contract, and he is more likely to win in events when he makes a low bid, hence in events where his prediction is low. The consequence is that conditional on winning, prediction errors do not cancel out anymore, because in these events where bidder i's prediction is low, bidder i also tends to be optimistic about the cost C_i (see Lemma 1 below). In other words, the auction induces **a selection bias** in favor of optimistic bidders.

Formally, we will prove the following Proposition, which says that unless he wins with probability one, bidder i is subject to the winner's curse.

Proposition 1. *Under Assumption* 1*, if* $1 > \Pr\{i \textit{ wins}\} > 0$*, then* $\Delta_i > 0$*.*

Before proving Proposition 1, we show the following result, which says that conditional on the event where bidder i makes a low prediction of cost (say, below a threshold p), bidder i tends to be optimistic about his cost. Proposition 1 will then follow because bidder i only wins in events where is prediction is below $p = \min_{j \neq i} \widehat{Y}_j$, the smallest prediction made by the other players.

Lemma 1. *Consider p such that $0 < \Pr\{\widehat{Y}_i < p\} < 1$. Then, under Assumption 1, we have: $E[H_i \mid \widehat{Y}_i < p] > 0$.*

Intuitively, for any cost realization, low realizations of the prediction $\widehat{Y}_i$ coincide with low realizations of the error term $\widetilde{\varepsilon}_i$, hence the optimism.

Proof. Consider the random variable $Z_i = p - C_i$. Under Assumption 1, $\widehat{Y}_i = X_i = C_i + \widetilde{\varepsilon}_i$ and $E[H_i \mid \widehat{Y}_i] = E[C_i - \widehat{Y}_i \mid \widehat{Y}_i]$, hence we have

$$E[H_i \mid \widehat{Y}_i < p] = E[-\widetilde{\varepsilon}_i \mid \widetilde{\varepsilon}_i \leq Z_i].$$

Since $E\widetilde{\varepsilon}_i = 0$, and since ε_i is independent of costs, for any realization $z_i \in Z_i$ that falls within the support of $\widetilde{\varepsilon}_i$, we have[10]

$$E[\widetilde{\varepsilon}_i \mid \widetilde{\varepsilon}_i \leq Z_i, Z_i = z_i] < 0. \tag{3.1}$$

Since $\Pr\{\widetilde{\varepsilon}_i \leq Z_i\} \in (0,1)$, the supports of $\widetilde{\varepsilon}_i$ and Z_i must overlap,[11] hence we obtain the desired inequality. ■

Proposition 1 is then obtained as an immediate corollary of Lemma 1:

Proof of Proposition 1: Define $P = \min_{j \neq i} \widehat{Y}_j$. We have:

$$\Delta_i = E[H_i \mid \widehat{Y}_i < P] \tag{3.2}$$

Since $1 > \Pr\{i \text{ wins}\} > 0$, the support of P and $\widehat{Y}_i$ must overlap, and the result then follows from Lemma 1.[12] ■

Comment 1: Proposition 1 is closely connected to Capen and al.(1971)'s insight: their result would correspond to the case where the distributions over costs are degenerate and concentrated on the same value for all bidders. Proposition 1 makes clear that their insight is valid more generally, and that it does not rely on costs across bidders being common or interdependent. In Capen and al., costs are common, but this

[10] If the error term was not independent of cost, inequality (3.1) would still hold under the assumption that conditional on cost, the error term is unbiased, i.e. $E[\widetilde{\varepsilon}_i \mid C_i = c] = 0$ for all cost realizations c.

[11] This is because each support is an interval and because $\widetilde{\varepsilon}_i$ and Z_i both admit a density that is everywhere positive on its support.

[12] For realizations p of P that fall above the support of $\widehat{Y}_i$, $E[H_i \mid \widehat{Y}_i < p] = EH_i = 0$.

is a mere consequence of their (implicit) assumption that the distribution over costs is degenerate, combined with a symmetry assumption.

Comment 2: While Proposition 1 illustrates that *competition* induces a selection bias in favor of optimistic bidders, Lemma 1 illustrates that a similar selection bias may also obtain without competition, when contractor i is just given an option to contract at a some price p (chosen as in Lemma 1). Assume that contractor i accepts the contract whenever he expects positive profits, that is whenever $p - \widehat{Y}_i$ is positive. Then the lower the error term, the more likely he accepts the contract, and as a result, conditional on accepting the contract, contractor i is too optimistic about his cost.[13]

One application of this observation concerns the winner's curse in buyer-seller relationships.[14] If, as in Akerlof (1970), valuations of the buyer and the seller are interdependent, and if the buyer fails to realize this, he will be disappointed ex post by the value of the object he bought. One corollary of Lemma 1 is that the same phenomenon may arise when valuations are drawn from independent distributions and the buyer only gets an imperfect estimate of his valuation. If he makes prediction errors, he will be more likely to buy whenever his prediction is optimistic.

Comment 3: A recent insight due to van den Steen (2004) is that relative overoptimism (i.e. the fact that an agent tends to have an optimistic perception of his own prospect, relative to the way others view his own prospect) may stem from the combinations of two things: the fact that the agent makes estimation errors in evaluating various alternatives, and the fact that there are various alternatives to choose from. In a similar vein, a corollary of Proposition 1 is that the same combination of factors (estimation errors and choice among various alternatives) generates optimism (not relative to other's perceptions, but relative to the true prospects). Indeed, assume there is a single agent, and interpret the index i as one of the possible projects that the agent may undertake. Choosing the project for which he has the lowest cost estimate is equivalent

[13] Lemma 1 is closely connected to Brown (1974)'s insight: When a firm makes estimation errors in evaluating the value of a project, and undertakes it whenever the estimation is above a threshold, the evaluation ends up being too optimistic on average. As in Capen and al., the analysis of Brown corresponds to the case where the distribution of cost would be degenerate.

[14] This winner's curse in bilateral negotiations has been examined in experiments by Samuelson and Bazerman (1985).

to selecting the agent's which has the lowest cost estimate in the auction. Proposition 1 then says that whichever project the agent ends up selecting (optimally), his estimation on that project will be too optimistic on average.

Comment 4: Proposition 1 easily generalizes to more general specifications of the prediction function. In particular, it holds under the "computation error specification" mentioned earlier.[15] It also hold whenever the two following properties are satisfied:

(i) $EH_i = 0$, which means that on average prediction errors cancel out, and

(ii) $E[H_i \mid \widehat{Y}_i < p]$ is decreasing in p, which means that lower predictions corresponds to higher levels of optimism.[16]

3.2. Implications for profits.

We investigate the implication for profits in the symmetric case. When there are few bidders, the effect on profit is not dramatic because the winner receives a payment $P = \min_{j \neq i} \widehat{Y}_j$ which may be substantially higher than his own prediction ($\widehat{Y}_i$). When competition gets fierce however, the difference between the prediction and the payment received vanishes, and because of the curse identified above, the winner eventually makes losses. Formally, we have:

Proposition 2. *Consider the symmetric case. If there are only two contractors competing, expected profits are positive. With sufficiently many competitors, expected profits become negative.*

Under Assumption 1, $P = \min_{j \neq i} X_j$, and expected profits can be written as

$$\Pi_i = \Pr\{X_i < P\}[E[P - X_i \mid X_i < P] - \Delta_i].$$

[15] Under the computation error specification, we have $E[H_i \mid \widehat{Y}_i < p] = E[-\eta_i \mid \eta_i + Y_i^* < p]$, which again is positive since η_i is an independent variable with zero mean.

[16] These conditions for example hold for the case examined in footnote 7. Indeed, we have:

$$E[H_i \mid \widehat{Y}_i] = -\frac{\mu(\lambda) - \mu(1)}{\mu(\lambda)}(\widehat{Y}_i - c_0)$$

So when $\lambda < 1$, contractors put more weight than they should on the estimate X_i, and as a result,$E[H_i \mid \widehat{Y}_i]$ is decreasing in $\widehat{Y}_i$ (hence a fortiori, $E[H_i \mid \widehat{Y}_i < p]$ is decreasing in p).

When competition gets fierce, the term $E[P - X_i \mid X_i < P]$ vanishes, so profits become negative because Δ_i is positive and does not vanish.

With only two competitors, contractor 1 wins when $C_2 + \varepsilon_2 > C_1 + \varepsilon_1$ and then makes a profit equal to $C_2 + \varepsilon_2 - C_1$. So Π_1 can be re-written as:

$$\Pi_1 = \Pr\{1 \; wins\}[E[C_2 - C_1 \mid C_2 - C_1 > \varepsilon_1 - \varepsilon_2] + E[\varepsilon_2 \mid \varepsilon_2 > C_1 + \varepsilon_1 - C_2].$$

Since $E\varepsilon_2 = 0$ and since $EC_2 = EC_1$, both terms on the right-hand side are positive, hence Π_1 is positive as well. ■

3.3. Comparative statics.

How does the curse change when competition increases? When competition increases, low realizations of $P = \min_{j \neq i} \widehat{Y}_j$ are more likely, hence when bidder i wins, he tends to have a lower prediction $\widehat{Y}_i$. If a lower prediction implies stronger optimism, then more competition should imply higher optimism conditional on winning, hence a stronger curse. Indeed, we assume below that $E[H_i \mid \widehat{Y}_i]$ is decreasing in $\widehat{Y}_i$, which means that bidder i is more optimistic when he has a lower prediction. We have:

Proposition 3. *Assume $E[H_i \mid \widehat{Y}_i]$ is decreasing in $\widehat{Y}_i$. Then Δ_i increases with the number of bidders.*

The proof is standard and relegated to the Appendix.

A similar conclusion holds for the effect of competition on welfare. If the correct prediction Y_i^* is smaller for lower values of the estimate X_i, or more generally, if $E[Y_i^* \mid \widehat{Y}_i]$ is increasing in $\widehat{Y}_i$, then competition will increase welfare, despite the fact that the winner gets more optimistic.

It is worth pointing out however that this *monotonicity is not guaranteed* in general. To see why, consider a case where errors are correlated with costs and where estimation errors are more pronounced for higher realizations of cost. Then more competition may make more likely the selection of contractors who simultaneously get a large realization of cost and a large (but negative) error term.[17] In that case, competition may be detrimental to welfare.[18]

4. Discussion.

We have so far focused on the case where costs and errors are drawn from independent distributions, and we have found that Capen and al.'s explanation for the winner's curse, based on the selection of optimistic bidders, applied to that setting. This is in contrast with the standard "informational" explanation: when costs and estimates are drawn from independent distributions, other bidders' bids or estimates cannot help a bidder improve his assessment of own cost; hence failing to take into account that information cannot generate a curse for winners.

In this Section, we wish to pursue our comparison between the two explanations and deal with the case of common costs, that is, the case where the individual cost C_i is the same for all bidders. We shall denote by C the random variable defining this common cost [19] and focus on the symmetric case.

We start with the standard "informational" explanation, based on the failure of bidders to recognize that other bidders' bids convey information about the cost.

[17] Such a correlation could arise quite naturally. If completion of the contract requires various skills, those who master them (hence are presumably more cost efficient) should also be better at estimating the cost of completing the contract.

[18] A simple example that illustrates this point is the following. Assume costs are i.i.d. on $[1, 2]$, and that $X_i = C_i + k\zeta_i(C_i - 1)$, with ζ_i uniform on $[-1, 1]$, and $k \in (0, 2)$. [So the error term is equal to $k\zeta_i(C_i - 1)$, it is correlated with cost, but since $E[k\zeta_i(C_i - 1) \mid C_i] = 0$, Proposition 1 applies - see footnote 10].

When $k < 1$, the support of X_i is $[1, 2 + k]$, and as competition increase, winners tend to have a realization of X_i closer to 1, which means higher welfare and vanishing optimism (so this is an example of a case in which Proposition 3 does not apply).

When $k > 1$ however, the support of X_i is $[2 - k, 2 + k]$, and as competition increase, winners tend to have a realization of X_i closer to $2 - k$, which means highest possible cost realization, and highest possible optimism.

[19] For any realization c of C, we have $C_i = c$ for all i.

4.1. Case (a): The standard "informational" explanation.

Following the popular explanation for the winner's curse, we first consider the case (labelled case a) where each bidder i is assumed to correctly compute $Y_i^*(X_i) = E[C_i \mid X_i]$, but to fail to recognize that other bidders' estimates convey information about his own cost (for example because he erroneously believes that costs are distributed independently across bidders.[20]). Bidder i's prediction is thus

$$\widehat{Y}_i^a(X_i) = Y_i^*(X_i),$$

and it is independent of whether he wins or not. Define Δ_i^a as the expected difference between realized and predicted costs, conditional on winning. That is, since we consider the symmetric case,

$$\Delta_i^a \equiv E[C - \widehat{Y}_i^a \mid X_i \leq X_j \text{ for all } j].$$

For any realization x of X_i, we have:

$$\begin{aligned} E\left[C - \widehat{Y}_i^a \mid X_i = x, X_i \leq X_j \text{ for all } j\right] &= E\left[C \mid X_i = x, X_i \leq X_j \text{ for all } j\right] - E\left[C \mid X_i = x\right] \\ &> 0, \end{aligned}$$

implying that

$$\Delta_i^a > 0.$$

This inequality captures the standard "informational" explanation for the winner's curse, and Δ_i^a measures the size of the curse for bidder i. Δ_i^a is positive (hence bidder i is too optimistic about his cost) because bidder i ignores the information conveyed by winning. Unlike in the previous section however, optimism does not stem from a selection bias in favor of optimistic bidders. If bidder i was not competing with other bidders, his prediction $\widehat{Y}_i^a(X_i)$ would be neither optimistic nor pessimistic. It would be the correct one, given the information that he possesses. His prediction $\widehat{Y}_i^a(X_i)$ *becomes* optimistic when he competes with other bidders and wins, because he (erro-

[20] This formalization corresponds to the fully cursed equilibrium examined in Eyster Rabin (2002).

neously) believes that winning conveys no information (and because winning is actually bad news).

4.2. Case (b): The explanation based on selection of optimistic bidders.

Following the analysis of Section 3, we now consider the case (labelled case b) where bidders are assumed to take their estimate at face value, and make the following prediction:

$$\widehat{Y}_i^b(X_i) = X_i.$$

We also assume that this prediction is *independent* of whether he wins or not, either because he has no doubt about his own cost being equal to X_i (extreme overconfidence), or because, as in case (a), bidders erroneously believe that costs are drawn from independent distributions. We will return to this assumption shortly (see subsection 4.4), and discuss cases where bidders are less confident in their estimate and may be aware that costs are identical across bidders (and thus attempt to draw inferences from winning).

On average, bidders' prediction coincides with that of case (a):

$$E\widehat{Y}_i^b = E\widehat{Y}_i^a$$

The size of the curse however differs. *Not only* does bidder i ignore the information conveyed by winning, *but he also* makes additional errors in predicting costs. These errors exacerbate the curse, for the reason identified in Section 3. Defining

$$\Delta_i^b \equiv E[C - \widehat{Y}_i^b \mid X_i \leq X_j \text{ for all } j],$$

we have:

$$\Delta_i^b - \Delta_i^a = E[Y_i^*(X_i) - X_i \mid X_i \leq X_j \text{ for all } j],$$

which by Lemma 1 is positive. The intuition is identical to that given before. Compared to case (a), bidder i may be too optimistic about his cost ($\widehat{Y}^b < Y_i^*$) or too pessimistic ($\widehat{Y}^b > Y_i^*$). Because competition tends to select optimistic bidders, winners are even more optimistic than under case (a), and $\Delta_i^b - \Delta_i^a$ captures the magnitude of that effect.

In summary, case (b) combines the two effects: the failure of bidders to recognize that winning conveys information on cost (i.e. the "informational effect"), and the endogenous selection of the more optimistic bidders (i.e. the "selection effect").

4.3. Illustration.

To illustrate, we compare the magnitude of each effect in a simple example. Assume that

$$C = c^0 + \gamma \text{ and } X_i = C + \widetilde{\varepsilon}_i$$

with γ and $\widetilde{\varepsilon}_i$ drawn from independent distribution. Then we have

$$\Delta_i^b = E[-\widetilde{\varepsilon}_i \mid \widetilde{\varepsilon}_i < \widetilde{\varepsilon}_j \ \forall j] \text{ and } \Delta_i^a = E[C - Y_i^*(X_i) \mid \widetilde{\varepsilon}_i < \widetilde{\varepsilon}_j \ \forall j]$$

We compare Δ_i^a and Δ_i^b under two distinct assumptions:

(i) We start with the case where γ and $\widetilde{\varepsilon}_i$ are drawn from identical distributions. Then it is easy to check that:[21]

$$\Delta_i^b - \Delta_i^a = \Delta_i^a.$$

Hence in that case, the informational effect and the selection effect both have the same magnitude.

(ii) We now turn to the case where γ is concentrated around 0. [22] A bidder who would know the distribution over cost and errors would realize that costs remain close

[21] To see why, observe that since γ and $\widetilde{\varepsilon}_i$ are drawn from identical distributions, $E[\gamma \mid X_i] = E[\widetilde{\varepsilon}_i \mid X_i]$, so

$$Y_i^*(X_i) = c^0 + E[\gamma \mid X_i] = c^0/2 + [c^0 + E[\gamma + \widetilde{\varepsilon}_i \mid X_i]]/2 = c^0/2 + X_i/2$$

hence, since γ and $\widetilde{\varepsilon}_i$ are drawn from independent distributions,

$$E[C - Y_i^*(X_i) \mid \widetilde{\varepsilon}_i < \widetilde{\varepsilon}_j \ \forall j] = c^0 - c^0/2 - c^0/2 + E[-\widetilde{\varepsilon}_i/2 \mid \widetilde{\varepsilon}_i < \widetilde{\varepsilon}_j \ \forall j]$$

implying that

$$\Delta_i^a = \Delta_i^b/2$$

[22] This is the case examined by Capen and al.

to c^0, whatever signal X_i is observed. Thus $Y_i^*(X_i) \simeq c^0$, and

$$\Delta_i^a \simeq 0.$$

In contrast, Δ_i^b depends only on the distribution of errors, hence it remains unchanged (and significant) even as γ gets concentrated around 0. Thus in that case, the winner's curse derives exclusively from the selection bias, and not from bidders ignoring information embodied in others' bids.

4.4. Over-confidence and the winner's curse: further comments.

Throughout the paper, we have paid particular attention to the case where bidder i makes prediction errors concerning *his own cost*. This corresponds to situations where bidders understand that their estimation is subject to errors, but downplay the magnitude of their own errors; the extreme case being one where bidders believe that they are not making any estimation errors. We have found that this assumption alone could be responsible for a winner's curse in the private cost setting.

When we moved to the common cost setting, we have added the assumption that bidders fail to recognize that they are facing a common value setting and that winning means that they have the lowest cost estimate. Under this additional assumption, we have found that whether bidders downplay the magnitude of their errors (case (b)) or not (case (a)), they fall prey to the winners' curse; and we could distinguish between the standard *"informational effect"*, and what we called, the *"selection effect"*.

In many settings however, bidders presumably do realize that they are facing a common value setting. This is certainly the case for example in the class room experiment examined by Bazerman and Samuelson (1983), where students are asked to bid for a jar filled with coins. In such settings, one key (empirical) issue is whether bidders would modify their prediction after being told, say, the distribution over other bidders' predictions.[23] It would be quite reasonable to expect bidders to be influenced by others' predictions, hence (unless they do not realize the connection between winning and the

[23] This is issue could be easily adressed in an experiment, but to our knowledge it has not been adressed yet.

fact that others had a smaller prediction) to *adjust* their prediction of costs conditional on winning. If the adjustment is strong, we should expect the winner's curse to be reduced substantially; and otherwise to persist.

What can be said concerning the magnitude of this adjustment?

First, in the extreme case where bidders believe that they are not making any errors, information concerning other bidders' estimates will not be used, hence the winner's curse will persist. More generally, the more bidders are *over-confident* in own signals (i.e. the more they downplay the magnitude of their own errors), the more limited the adjustment will be, and this should be true even if players have a correct idea of the joint distribution over other players' estimation errors.[24]

So over-confidence in own signals may be responsible for a winner's curse in the common cost setting as well, *even when* players do realize that they are facing a common value setting, and *even when* they realize that winning conveys information on others' estimates. But the winner's curse now obtains through *two* possible channels: the "selection effect", because over-confidence in own signals induces bidders to overweigh own estimate X_i; and the "informational effect", because over-confidence in own signals also leads bidders to underweigh others' estimates.

Of course, other reasons (besides over-confidence in own signals) may lead bidders to disregard or put too little weight on others' estimates: even when players have a correct idea of the distribution of their own costs and estimations, and correctly compute conditional expectations, they may have an inflated idea of the magnitude of errors made by others or of the correlation of these errors,[25] with the consequence that they fail to incorporate in their prediction as much information as they should concerning other

[24] For example, in the case of normal distributions examined in footnote 3, if bidder i believe that $X_i = C + \lambda\varepsilon_i$ and $X_j = C + \varepsilon_j$, we have

$$E[C \mid X_i, X_j] = \mu_i(\lambda)X_i + \mu_j(\lambda)X_j + (1 - \mu_i(\lambda) - \mu_j(\lambda))c_0$$

with $\mu_i(\lambda) = \frac{\nu^2}{(\lambda\nu)^2+\nu^2+(\lambda\nu^2/\sigma)^2}$ and $\mu_j(\lambda) = \frac{(\lambda\nu)^2}{(\lambda\nu)^2+\nu^2+(\lambda\nu^2/\sigma)^2}$. So the smaller λ, the more weight contractor i puts on own signal X_i, and the lesser weight contractor i puts on contractor j's signal X_j.

[25] The experiment conducted by Bazerman and Samuelson (1983) actually reveals that errors may be highly correlated and biased: the average prediction of students were 37% below the true value of the jar.

bidders' estimates.

Which type of error (downplaying own estimation errors or inflating others') is more prevalent in practice is an empirical question that would be worthwhile investigating in the lab.

5. Conclusion

This paper has attempted to clarify the difference between two types of arguments used in the literature to account for the winner's curse. Both arguments are based on mistakes made by the bidders. The first one is based on the failure of bidders to (fully) incorporate in their predictions the information on others' valuations conveyed by winning. The second one is based on bidders not realizing that (or the extent to which) their value estimate is subject to errors along with the fact that they are more likely to win when their estimate is optimistic.

One implication of our analysis is that the winner's curse phenomenon may be more *widespread* and *persistent* than generally thought.

Indeed, the second argument applies very broadly, including to settings to which the first argument does not apply (e.g. to settings where valuations are drawn from independent distributions, or to common value settings where bidders share the same data.[26,27]). In addition, a single and very common cognitive bias (i.e. over-confidence in the accuracy of own signal) may be responsible for both mistakes. Finally, over-confident bidders will continue to fall prey to the winner's curse even if they are taught that winning conveys information on others' estimates.

Several issues remain to be addressed. (i) In theory, various types of errors may explain the winner's curse: failure to recognize that winning conveys information about others' estimate, over-confidence in own signals, under-confidence in the accuracy of other's signals, and testing the relevance of each explanation deserves further research.

[26] In these settings, other bidders' bids convey no additional information (because there are no additional information to be obtained), hence the first argument does not apply.

[27] One good illustration is the class room example mentioned earlier where students are asked to bid for a jar filled with coins (See Bazerman and Samuleson 1983). They all share the same data (a view of the jar), and yet they fall prey to winner's curse.

(ii) We have portrayed the behavior of somewhat naive bidders who often take their estimates at face value and bid accordingly, and never learn that their choice is suboptimal. In the long run, as experience builds up, we might expect bidders to become aware that their choice is suboptimal (in particular when these choices lead to losses) and adjust their bids accordingly. This is a line of research we have started in Compte (2001).[28]

Appendix

Proof of Proposition 3: Let φ_i denote the distribution over X_i induced by f_i and g_i, and $Q_i(x) = \Pr\{X_i > x\}$. Also let $h_i(x) = E(H_i \mid \widehat{Y}_i = x)$. With n bidders, we have:

$$\Delta_i^{(n)} = \frac{\int h_i(x)\psi_i(x)dx}{\int \psi_i(x)dx}$$

where $\psi_i(x) = \varphi_i(x) \prod_{j=1, j \neq i}^{n} Q_j(x)$, and where the superscript (n) indicates the number of bidders.

With one additional bidder, say bidder $n+1$, the bias Δ_i becomes

$$\Delta_i^{(n+1)} = \frac{\int h_i(x)Q_{n+1}(x)\psi_i(x)dx}{\int Q_{n+1}(x)\psi_i(x)dx}$$

Since both h_i and Q_{n+1} are decreasing functions, and since the expectation of the product of two decreasing functions is larger than the product of the expectations, we have:

$$\frac{\int h_i(x)Q_{n+1}(x)\psi_i(x)dx}{\int \psi_i(x)dx} \geq \frac{\int h_i(x)\psi_i(x)dx}{\int \psi_i(x)dx} \times \frac{\int Q_{n+1}(x)\psi_i(x)dx}{\int \psi_i(x)dx}$$

which implies , $\Delta_i^{(n+1)} / \Delta_i^{(n)} \geq 1$. ■

[28] In Compte (2001), we propose a model along this line, where bidders learn to set (optimally) a uniform mark-up on their cost estimate. In this model we find that private value *and* common value settings both yield similar qualitative predictions concerning the effect of competition on bidding behavior, namely, increased cautiousness.

References

[1] George Akerlof, “The market for ‘lemons’: Quality uncertainty and the market mechanism”. *The Quarterly Journal of Economics*, August 1970, 488–500.

[2] Bazerman, Max and Samuelson, William “I won the auction but don’t want the price”, *Journal of Conflict Resolution*, December 1983, 27; pp 618-34

[3] Keith Brown, “A note on the Apparent Bias of Net Revenue Estimates for Capital Investment Projects”, *Journal of Finance*, September 1974, 29, pp 1215-1216

[4] E.C. Capen, R.B. Clapp, and W. M. Campbell, “Competitive Bidding in High Risk Situations,” *Journal of Petroleum Technology*, June 1971, 23, pp 641-53.

[5] Olivier Compte, “The winner’s curse in a private value setting”, June 2001, mimeo, CERAS (http://www.enpc.fr/ceras/compte/winner.pdf)

[6] Kent Daniel, David Hirshleifer, and Avanidhar Subrahmanyam, “Investor psychology and security market under- and over-reactions”, *Journal of Finance* 1998, 53, 1839-1885.

[7] J.B. De Long, A. Shleifer, L.H. Summers, and R.J. Waldmann, “The Survival of Noise Traders in Financial Markets,” *Journal of Business*, 1991, 64, 1-19.

[8] Erik Eyster and Matthew Rabin, ”Cursed Equilibrium”, August 2002, Berkeley Working Paper E02-320.

[9] Jacob K. Goeree and Theo Offerman, “Winner’s Curse Without Overbidding”, 2002, Working paper

[10] John H. Kagel and Dan Levin, “The Winner’s Curse and Public Information in Common Value Auctions” *American Economic Review* December 1986, 76 (5) pp 894-920.

[11] John H. Kagel and Dan Levin, “Bidding in Common value auctions: a Survey of Experimental Research”, in J.H. Kagel and D. Levin, *Common Value Auctions and the Winner’s Curse*, Princeton: Princeton University Press. 2002

[12] A. Kyle and F.A. Wang, "Speculation Duopoly with Agreement to Disagree: Can overconfidence survive the market test?" *Journal of Finance*, 1997, 52, 2073–2090.

[13] Paul Milgrom, "Auctions and Bidding: A Primer", *The Journal of Economic Perspective*, Summer 1989, 3 (3), pp 3-22.

[14] Terrence Odean, "Volume, Volatility, Price, and Profit when all Traders are Above Average," *Journal of Finance*, 1998, 53, 1887–1934.

[15] Samuelson, William and Bazerman, Max "The Winner's Curse in Bilateral Negotiations," in V. L. Smith (ed.) *Research in Experimental Economics*, Vol. 3, Greenwich, Conn.: JAI Press, 105-137, 1985

[16] Eric van den Steen, "Rational Overoptimism (and Other Biases)", *American Economic Review*, 2004, forthcoming.

[17] Richard Thaler. The winner's curse. New York: Free Press, 1992