Australian School of Business
Working Paper

Never Stand Still   Australian School of Business

Australian School of Business Research Paper No. 2013 ECON 28

The Satisficer's Curse

Robert E. Marks

The Satisficer's Curse

Robert E. Marks
School of Economics
University of New South Wales
Sydney, NSW 2052, Australia

bobm@agsm.edu.au

ABSTRACT: Following the Winner's Curse and the Optimizer's Curse, this paper introduces the Satisficer's Curse. The Winner's Curse requires competition between agents in an auction for, usually, a common-value item; the recently named Optimizer's Curse is a systematic overvaluation when the decision maker is choosing the highest valued prospect of a set of uncertain future outcomes. The Satisficer's Curse is a systematic overvaluation that occurs when any uncertain prospect is chosen because its estimate exceeds a positive threshold. It is the weakest version of the three curses, all of which can be seen as statistical artefacts.[1]

1. Introduction

Thirty-six years ago a trio of oil men (Capen et al. 1970) observed and named the phenomenon of the Winner's Curse in auctions, later epitomised by the saying, "The good news is you won; the bad news is you paid too much". Their point-estimate model of selection bias was generalised by Harrison and March (1984), who assumed Gaussian distributions in their modelling, and further generalised by Smith and Winkler (2006), who argued that there was no need for the explicit competition of the auction mechanism, and who coined the phrase, the "Optimizer's Curse," to argue that any decision that chooses the best prospect from a set of possibilities will fall heir to "post-decision" disappointment, on average.

The first paper to draw attention to the post-decision disappointment associated with internal capital investment decisions was Brown (1974), but Harrison and March (1984, p.27) also noted that when a decision-maker was choosing the best option among a predetermined number of alternatives, or when the alternatives were considered sequentially until one was identified as satisfying some predetermined aspiration level or hurdle, then post-decision disappointment would occur, on average. Compte (2004, p.9) also commented that selection bias would occur when choosing the best from a set of alternatives, without any recourse to psychological motivation (Camerer and Lovallo 1999, Lovallo and Kahnemann 2003, Tiwana et al. 2007). See Statman and Tyebjee (1985) and van den Steen (2004) for surveys of evaluation biases and their psychological and other foundations.

---

1. I wish to thank Dan Lovallo and Danny Kahnemann for the presentation which sparked this paper, even if they disagree with my thesis, and Joshua Gans, Frank Milne, Anna Gunnthorsdottir, Geoff Eagleson, Sharyn Roberts, and participants at the 25th Australasian Economic Theory Workshop, A.N.U., for their helpful comments.

This paper formalises the post-decision disappointment flowing from selection bias when the decision-maker is choosing an alternative that satisfies some (positive) hurdle in expected performance, the Satisficer's Curse.[2]

2. The Formal Model

Let S_i be the expected difference between the predicted value $\hat{Y}_i$ and the ex-post realised value v_i of an uncertain future prospect i, conditional on the predicted value exceeding a positive threshold $p > 0$. We write this as

$$S_i \equiv E[\hat{Y}_i - v_i | \hat{Y}_i > p > 0] \tag{1}$$

where the predicted value $\hat{Y}_i$ of prospect i is assumed to be the sum of the true value μ_i and an error ε_i:

$$\hat{Y}_i = \mu_i + \varepsilon_i \tag{2}$$

Choosing prospect i when its predicted value exceeds a theshold is "satisficing."[3] That is, the predicted value is "good enough" to choose or accept, rather than being the best, or optimized. Previous work by Smith and Winkler (2006) demonstrated that choosing the optimal prospect would result in disappointment, on average. They called this the "Optimizer's Curse." Here we weaken their findings, with the "Satisficer's Curse."

Theorem 1: Absent a conditional Bayesian expectation of value, S_i will be strictly positive.

Before proceeding to prove this, we need some further scaffolding.

In a Bayesian world, the joint distribution over the ex-post realised value v_i and the predicted value $\hat{Y}_i$ would be known. This would allow computation of the conditional Bayesian expectation of value, given a specific ex-post realisation v_i of the value estimate V_i:

$$Y_i^*(v_i) \equiv E[\mu_i | V_i = v_i].$$

Assumption 1: Decision makers assume that $\hat{Y}_i = V_i$, that is, that the predicted value equals the estimated value. This is equivalent to assuming (perhaps implicitly) that

---

2. Are we cursed with a plethora of "curses"? Apart from the Winner's Curse, there is the "news" curse (Goeree and Offerman 2003), the "loser's curse" (Massey and Thaler 2006), the Optimizer's Curse, and now the Satisficer's Curse.

3. Simon (1957) introduced the word "satisfice" as a description of non-optimizing decision making; satisficing is now institutionalized as a means of making multifarious decisions.

$E[\hat{Y}_i] = \mu_i$, which implies ignoring the error ε_i.[4]

Note that $E[\hat{Y}_i] = E[Y_i^*(v_i)]$, that is, unconditional errors cancel. We are, however, concerned with conditional expectations: the expected prediction error $(\hat{Y}_i - v_i)$ conditional on $\hat{Y}_i > p > 0$, the predicted value exceeding the positive threshold.

If Theorem 1 is correct, then for any ex-post value realisation v_i, high realisations of the predicted value $\hat{Y}_i$ coincide with high realisations of the error term ε_i, and hence coincide with over-valuation. That is, the expected difference between the predicted and realised values, given that the predicted value exceeds a positive threshold, is positive.

Definition 1: Prospect i exhibits the Satisficer's Curse when $S_i > 0$.

Definition 2: Define the decision-maker's *optimism* H_i as the difference between the predicted value $\hat{Y}_i$ and the conditional Bayesian predicted value $Y_i^*(v_i)$:

$$H_i(V_i, \hat{Y}_i) \equiv \hat{Y}_i - Y_i^*(v_i). \qquad (3)$$

It follows that $S_i = E[H_i | \hat{Y}_i > p > 0]$.

Proof (of Theorem 1, after Compte's (2004) Lemma 1): Assumption 1 implies that $\hat{Y}_i = V_i$, so the event $\{\hat{Y}_i > p > 0\}$, that is, that the event that the predicted value exceeds a positive threshold, is equivalent to the event $\{\varepsilon_i > Z_i\}$, the error ε_i exceeds a value Z_i, defined as the difference between the positive threshold p and the true value μ_i:

$$Z_i \equiv p - \mu_i \qquad (4)$$

Then, taking expectations of both sides of Eq. (3), the expected optimism, given that the predicted value exceeds the positive threshold, is given by

$$\begin{aligned} E(H_i | \hat{Y}_i > p) &= E(V_i - Y_i^*(v_i) | \varepsilon_i > Z_i) \\ &= E(\varepsilon_i | \varepsilon_i > Z_i) \end{aligned}$$

But, by construction $E(\varepsilon_i) = 0$, and ε_i is independent of values μ_i for any realisation $z_i \in Z_i$ that falls within the support of ε_i (which is unbounded if the p.d.f. g_i generating the error ε_i is Gaussian).

Thus we have $E(\varepsilon_i | \varepsilon_i > Z_i, Z_i = z_i) > 0$ if $Z_i > 0$.

Since Pr $\{\varepsilon_i \geq Z_i\} \in (0,1)$, the supports of ε_i and Z_i must overlap.

This follows because each support is an interval, and because ε_i and Z_i both admit of a density that is everywhere positive in its support, by definition. Hence

$$E[H_i | \hat{Y}_i > p > 0] > 0$$

4. Tiwana et al. (2007) might suggest that this is a version of their bounded rationality bias.

and Theorem 1 is proved. □

We can generalise this. If the predicted value $\hat{Y}_i$ is based on past observations, then it is not sufficient for future desired performance (or even future satisfactory performance) that an observed event has occurred in which the realised value v_i exceeded the threshold p, the positive hurdle. In general, a satisficing decison-maker will be disappointed in future: the *Satisficer's Curse.*

It is understood that setting performance hurdles does not guarantee better performance in the future, unless there is trending improvement of performance, so that the underlying stochastic process has a rising true value $\mu_{(t)}$. But the Satisficer's Curse is saying that, in general, future performance of a stationary stochastic process will not attain the positive threshold, at least in expected terms.

3. A Model of Project Selection

Assume n independent projects available for internal investment by a company, each of which has an independent value μ_i, described by a p.d.f. f_i, and an error term ε_i, described by a p.d.f. g_i. The estimate V_i of the ith project's value is the sum of the true value μ_i and the error term ε_i:

$$V_i = \mu_i + \varepsilon_i. \tag{5}$$

The error terms are assumed independent across projects, and are unbiased: $E[\varepsilon_i] = 0$. We assume that the distributions over values and errors are non-degenerate. Formally, we assume that the value μ_i and the error ε_i each admit of a density (as above, denoted by f_i and g_i respectively), that the support of each density is an interval, and that each density is positive in its support.

This model states that there is randomness in prediction (sampling), and randomness in the eventual project outcome. The first randomness is described by the error-term p.d.f. g_i, and the second by the p.d.f. f_i of the value term μ_i for project i. There is a single realisation v_i of the value estimate V_i, and (at most) a single realisation x_i of the random variable μ_i. Of course, if project j is not chosen, there is no realisation of its value.

In a Bayesian world the firm would know the joint distribution over the values μ_i and the estimates V_i. So the firm would be able to compute, for each realisation v_i of the estimate V_i, the conditional probability over the value, and hence the conditional Bayesian expectation of value denoted by $Y_i^*(v_i)$

$$Y_i^*(v_i) \equiv E[\mu_i \mid V_i = v_i] \tag{6}$$

Using the correct conditional Bayesian expectation of value would not guarantee the absence of post-decision disappointment ($Y_i^*(v_i) - x_i > 0$), since there would still be variance in the estimates, but on average there would be no overestimation of values: estimate V_i would be unbiased.

Here, in contrast, we assume that the firm is aware that values μ_i are distributed independently, but that, conditional on V_i, the firm forms an erroneous prediction $\hat{Y}_i$ of the value μ_i. (Thus $\hat{Y}_i \neq Y_i^*$ necessarily. This could also be explained by use of an erroneous conditional distribution of value μ_i given estimate V_i; other possible reasons are discussed below.)

Assumption 1 above states that the predicted value $\hat{Y}_i$ is given by the estimated value V_i. This assumption models the polar case in which the firm believes that the estimate V_i has a higher predictive content than it really has, by ignoring the error term in equation (5). This does not imply that $\hat{Y}_i$ is greater than Y_i^*, only that the firm ignores the error in estimation, and underestimates the error in the realisation x_i of μ_i.

To summarise: For some projects the firm will be too optimistic about the value, meaning that:

$$\hat{Y}_i > Y_i^*(v_i)$$

and for other projects the firm will be too pessimistic:

$$\hat{Y}_i < Y_i^*(v_i).$$

4. The Optimizer's Curse

We denote the difference between the actual prediction and the conditional Bayesian prediction from Equation (6) by the optimism H_i associated with project i: $H_i(V_i, \hat{Y}_i) \equiv \hat{Y}_i - Y_i^*(v_i)$, Note that under Assumption 1, by assumption $E(V_i) = E(\mu_i)$ (estimate V_i is an unbiased estimate of value μ_i), so that on average there is no optimism associated with any project i, and the prediction errors cancel:

$$E(H_i) = 0.$$

Let Δ_i denote the expected difference between predicted and realised values, conditional on being chosen, that is:[5]

$$\Delta_i \equiv E[\hat{Y}_i - \mu_i \mid i \text{ is chosen}] = E[\hat{Y}_i] - v_i \tag{7}$$

Definition 3: Project i exhibits the Optimizer's Curse when $\Delta_i > 0$.

In other words, the Optimizer's Curse refers to situations in which the value of the chosen project was overestimated. Following Compte (2004), note that Δ_i can be rewritten as:

$$\Delta_i = E[\hat{Y}_i - Y_i^* \mid i \;\text{ is chosen}] = E[H_i \mid i \;\text{ is chosen}], \tag{8}$$

since $E[Y_i^* \mid i \;\text{ is chosen}] = E[\mu_i \mid i \;\text{ is chosen}]$ by construction.

When we consider competing bidders in an auction, in the limiting case of a single bidder (who is thus certain to win), the prediction errors cancel, and the biider does not suffer

5. The following formal model is an adaptation from Compte (2004), who however treated the Winner's Curse in auction selection, not the Satisficer's Curse, a much weaker concept.

the Winner's Curse on average. But in the case of a firm choosing a single project from a set, with possibly erroneous predictions of the values of each, almost always the firm will suffer the Optimizer's Curse, even though each prediction is an unbiased estimate of that project's value. Only if there were a single project to choose from would the firm not experience the Optimizer's Curse on average.

How is this so? The act of choosing the project with the highest predicted (net) value induces a *selection bias* in favour of projects with (overly) optimistic value predictions.

Under what circumstances would such a firm *not* suffer a once-off occurrence of post-decision disappointment? When both of the following conditions are met:

a. when project k is chosen, because $\hat{Y}_k > \hat{Y}_i$ for all $i \neq k$, or $\hat{Y}_k > \max_{i \neq k} \hat{Y}_i$, *and*

b. when the highest value prediction $\hat{Y}_k$ (of project k) is less than the realisation v_k, so that $H_k \equiv \hat{Y}_k - Y_k^*(v_k) < 0$.

That is, $\Delta_k = E(\hat{Y}_k - \mu_k \mid k \text{ is chosen }) < 0$. Of course, that a single occurrence is profitable does not preclude the Optimizer's Curse from occurring over several repetitions: given the stochastic nature of the net returns, it is the expectation of these returns that indicates the existence of the Optimizer's Curse, or not.

Theorem 2: (From Compte's (2004) Proposition 1) With $\hat{Y}_i = V_i$ (Assumption 1), if $0 < \Pr\{i \text{ wins}\} < 1$, then $\Delta_i > 0$, that is, if the decision maker uses the naive forecast (Assumption 1), and the project could be chosen (its choice is neither certain nor impossible), then the project exhibits the Optimizer's Curse.

Theorem 2 will follow from Theorem 1 because project i is only chosen (wins) in events where its predicted value is equal to or greater than $\overline{p} = \max_{j \neq i} \hat{Y}_j$, the highest prediction across other projects.

Proof of Theorem 2: Define

$$\overline{p} \equiv \max_{j \neq i} \hat{Y}_i$$

We have

$$\Delta_i = E[H_i \mid \hat{Y}_i > \overline{p}],$$

given that i is preferred to all others. Since $0 < \Pr\{i \text{ is chosen}\} < 1$, the support of $\overline{p}$ and $\hat{Y}_i$ must overlap, for the same reason as above. Thus, from Theorem 1, the result follows:

$$\Delta_i > 0,$$

and Theorem 2 is proved (under Assumption 1): the chosen project exhibits the Optimizer's Curse. □

5. Discussion

Compte (2004) notes that our Theorem 2 is connected to Capen et al.'s (1971) insight, and relies on neither point estimates, nor values being common or interdependent. Its proof does not rely on Gaussian distributions, either. He further notes that Theorem 2 illustrates how competition induces a selection bias in favour of overly attractive projects. Theorem 1 illustrates that a similar selection bias may occur without competition, when a project is undertaken if it exceeds some positive hurdle p. If a project i is chosen whenever it looks attractive (whenever its NPV is greater than some positive threshold p), that is, whenever $(\hat{Y}_i - p)$ is positive, then the higher the error term, the more likely the project is to be undertaken, and, as a result, conditional on accepting the contract, project i is overly attractive.

The Satisficer's Curse is similar to the Peter Principle (when no competition for promotion or tenure exists, just a performance hurdle), although Lazear (2004) points out that those decisions are special cases in the estimates (our V_i) are based on past performance alone. There is no such restriction on how our estimates are derived; indeed, for many proposals there will be no past performance to observe.

Van der Steen (2004) argues in effect that overly attractive predictions may stem from, first, estimation errors, as discussed above, and, second, from the range of attractivenesses (i.e. diversity across projects).

A corollary of Theorem 2 is that the same combination of factors (viz. estimation errors and choice among various alternative projects) generates over-attractiveness (relative to true prospects). Selecting the project which appears to have the highest value (estimate) to the firm is equivalent to choosing the agent with the highest estimate of the item being sold in an auction. Theorem 2 says that whichever projects the firm ends up selecting (optimally) will turn out to have been valued optimistically, on average.

Will competition among firms reduce the Satisficer's Curse? As Massey and Thaler (2006) note, the Winner's Curse can persist in competitive markets because there are limits to arbitrage: the winners either go broke or learn; wiser heads must watch from the sidelines and hope for the former. "Since there is no way to sell the oil leases short, the smart money cannot actively drive the prices down." Since the Satisficer's Curse does not assume any interaction between firms, competition plays no direct role here.

6. Origins of Biases

As Compte (2004) suggests, we could have modelled Assumption 1 as:

$$V_i = \mu_i + \lambda \varepsilon_{i,} \tag{9}$$

where $\lambda \in [0,1)$. That is, the firm realises that estimation error is possible but downplays the magnitude of its own errors by a factor λ. The firm's prediction $\hat{Y}_i$ of the value μ_i is:

$$\hat{Y}_i \equiv E^{\lambda}[\mu_i \mid V_i],$$

where the superscript λ means that the expectation is taken assuming a joint distribution over V_i and μ_i is characterised by equation (9).

We have already shown that (Theorem 2) a project chosen under the naive Assumption 1 will exhibit the Optimizer's Curse (in its expected sense), and a moment's thought about equation (9) shows that only if the full error term ε_i is acknowledged ($\lambda = 1$) does $\hat{Y}_i = Y_i^*$; for any λ less than 1 Theorem 2 still holds.[6] Indeed, it is easily shown that Δ_i is decreasing in λ.

7. Learning to Avoid the Satisficer's Curse?

What about learning? Discussing sealed-bid tenders to sell in procurement auctions, Compte (2004) proposed a model in which bidders learn to set a mark-up on their cost estimates to reduce the risk of suffering the Winner's Curse, and argued that this leads to increased cautiousness in bidding, whether with private or common values.

In the case we consider of a firm choosing a prospect from a range of prospects, what is the decision-maker to learn? Should he or she ignore the ranking by predicted value because of the error terms ε_i? To do so would be to throw information away. Raising any return hurdle $\hat{p}$ that some projects are predicted to exceed will not obviate the Satisficer's Curse (from Theorem 1) so long as the error term is ignored ($\lambda = 0$) or discounted ($\lambda \in [0,1)$).

If the hurdle is an institutional threshold, then an understanding of the Satisficer's Curse should result in the institution learning to put procedures in place to reduce the prospect of performance reverting to the mean in future. For example, accreditation (of a business school to the AACSB, for instance) should be followed by the school using the accreditation inspection process to institutionalise assurance procedures for maintaining or even improving future performance, lest entropy increase after the hurdle has been surpassed and accreditation achieved, leading to consequent withdrawal of accreditation.

If such learning is for whatever reason not available to the decision makers or those who benefit from jumping the hurdle, then acknowledgement of the Satisficer's Curse should qualify expectations that future performance will reflect past estimates; on average it will not.

References

[1] Brown, K.C., A note on the apparent bias of net revenue estimates for capital investment projects, *Journal of Finance,* 29(4): 1215–1216, September 1974.

[2] Camerer, C. and D. Lovallo, Overconfidence and excess entry: an experimental approach, *American Economic Review,* 89: 306–318, 1999.

6. This is what Goeree and Offerman (2003) term the "news" curse: decision-makers neglect the fact that a high estimate makes a positive error more likely.

[3] Capen E.C., R.B Clapp, and V.M. Campbell, Competitive bidding in high risk situations, *Journal of Petroleum Technology,* 23: 641−653, June 1971.

[4] Compte, O., Prediction errors and the winner's curse, ENPC-CERAS working paper, September 2004. www.enpc.fr/ceras/compte/prediction2004.pdf (accessed 2006/05/15)

[5] Goeree J.K. and T. Offerman, Winner's Curse without overbidding, *European Economic Review,* 47(4): 625−644, 2003.

[6] Harrison J.R. and J.G. March, Decision making and postdecision surprises, *Administrative Science Quarterly,* 29: 26–42, 1984.

[7] Lazear, E.P., The Peter Principle: a theory of decline, *Journal of Political Economy,* 112(1 pt.2): S141−S163, 2004.

[8] Lovallo, D. and D. Kahnemann, Delusions of success: how optimism undermines executives' decisions, *Harvard Business Review,* 81(7): 56–, July 2003.

[9] Massey, C. and R.H. Thaler, The loser's curse: overconfidence vs. market efficiency in the National Football League draft, mimeo., March 2006. http://www.econ.berkeley.edu/~webfac/malmendier/e218_sp06/Thaler.pdf (accessed 2006/05/18)

[10] Simon, H.A., *Models of Man*, New York: John Wiley, 1957.

[11] Smith, J.E. and R.L. Winkler, The optimizer's curse: skepticism and postdecision surprise in decision analysis, *Management Science,* 52(3): 311−322, 2006.

[12] Statman, M. and T.T. Tyebjee, Optimistic capital budgeting forecasts: an experiment, *Financial Management,* pp. 27−33, Autumn, 1985.

[13] Tiwana, A., Wang, J., Keil, M., Ahluwalia, P., The Bounded Rationality Bias in Managerial Valuation of Real Options: Theory and Evidence from IT Projects, *Decision Sciences,* 38(1): 157−181, February 2007

[14] van der Steen, E. (2004) Rational overoptimism (and other biases), *American Economic Review,* 94(4): 1141−1151, Sept.