



Munich Personal RePEc Archive

Cost overruns and demand shortfalls – deception or selection?

Eliasson, Jonas and Fosgerau, Mogens

Centre for Transport Studies, KTH Royal Institute of Technology,
Stockholm, DTU Transport, Copenhagen

2013

Online at <https://mpa.ub.uni-muenchen.de/49744/>

MPRA Paper No. 49744, posted 11 Sep 2013 11:42 UTC

Cost overruns and demand shortfalls – deception or selection?

Jonas Eliasson^a, Mogens Fosgerau^{a,b}

^a Centre for Transport Studies, KTH Royal Institute of Technology, Stockholm

^b DTU Transport, Copenhagen

CTS Working Paper 2013:X

Abstract

A number of highly cited papers by Flyvbjerg and associates have shown that ex ante infrastructure appraisals tend to be overly optimistic. Ex post evaluations indicate a bias where investment costs are higher and benefits lower on average than predicted ex ante. These authors argue that the bias must be attributed to intentional misrepresentation by project developers. This paper shows that the bias may arise simply as a selection bias, without there being any bias at all in predictions ex ante, and that such a bias is bound to arise whenever ex ante predictions are related to the decisions whether to implement projects. Using a database of projects we present examples indicating that the selection bias may be substantial. The examples also indicate that benefit-cost ratios remains a useful selection criterion even when cost and benefits are highly uncertain, gainsaying the argument that such uncertainties render cost-benefit analyses useless.

Keywords: cost overruns, cost escalation, forecast accuracy, cost-benefit analysis, appraisal, selection bias, winner's curse.

JEL Codes: R40, R42.

1 INTRODUCTION

A large body of evidence shows that transport investments are often subject to cost overruns, and that costs have been underestimated on average (see e.g. van Wee (2007); a summary of several studies can be found in (Lundberg, Jenpanitsub, & Pyddoke, 2011)). This bias has been taken as a sign that cost overruns cannot be caused simply by “honest errors” in the ex-ante cost estimates. In particular, Bent Flyvbjerg and his associates have published a series of much cited papers (Flyvbjerg, 2008, 2009; Flyvbjerg, Holm, & Buhl, 2002; Flyvbjerg, Skamris Holm, & Buhl, 2004, 2005) that indicate a persistent bias in infrastructure project appraisals, where costs are systematically underestimated and benefits are systematically overestimated. They argue that the bias must be due to systematic misrepresentation by project promoters, and they use words such as deception and lie to describe what is going on. This is clearly a very serious critique. Random errors are explicitly rejected as an explanation for the observed forecast bias. In the words of Flyvbjerg (2009): *“If misleading forecasts were truly caused by technical inadequacies, simple mistakes, and inherent problems with predicting the future, we would expect a less biased distribution of errors in forecasts around zero.”* No supporting arguments for this claim are provided.

This paper will show, however, that such a bias can occur as a result of the selection process, without there being any bias at all in the forecasts ex ante. All it takes for bias to occur is that the selection of projects is related to ex ante predictions. Thus we show the Flyvbjerg argument to be invalid: It is perfectly possible that forecasts are actually unbiased, but that selection of the best projects, influenced by the same forecasts, leads to bias. It follows that it is not possible to conclude from the observation of ex post bias that that bias must be deliberate. Note that we do not claim that deliberate deception does not occur – in fact, there is substantial evidence that it does. Our point is that any selection process that is affected by predicted costs or benefits will yield biased outcomes. Hence, selection is a very plausible cause of observed biases in costs or benefits. Whether it is the *only* cause of biases in a given context is impossible to say without further evidence.

That selection processes from an unbiased population may yield biased outcomes is of course a general phenomenon. In the auction literature, a similar phenomenon is termed *winner’s curse* (Thaler, 1988). The name comes from sealed-bid, common-value auctions, where a number of bidders bid for an item of uncertain value, a value which is the same for all bidders. All bidders guess the true value of the item, place sealed bids according to these guesses, and the highest bidder gets to purchase the item. It is clear that, on average, the result will be that the winner ends up paying more than the value of the item – hence the name “winner’s curse”. In econometrics the same phenomenon is called selection bias (Heckman, 1979) and its presence means that the independent variables are rendered not independent of noise terms; this violates the basic statistical assumptions that are commonly made and more involved methods have to be used. A process similar to selection may also explain the apparent overconfidence exhibited when a majority of people rank themselves as better than average on easy tasks (e.g. driving) and worse than average on difficult tasks; for example, people who have not been involved in road accidents may rationally rate themselves as better than average drivers, this would not be overconfidence (Benoît & Dubra, 2011). “Regression to the mean” is a well-known statistical trap caused by selection: if participants in some experiment are selected based on some characteristic with random variation – say, having a high result on a test – then follow-up measurements of that characteristic (a repeated test, say) will show that the selected group has become more similar to a

control group (the test difference will be smaller). This trap may lead to false inferences concerning the effects of some kind of treatment.

Selection may cause systematic cost overruns and benefit shortfalls for transport investments. Imagine a decision maker faced with a number of alternative investments with uncertain cost estimates (ignore the uncertainty of benefits for the time being). Based on these estimates, the decision maker selects a number of the projects, with less costly projects being more likely to be selected. Now, even if the random errors of the initial cost estimates have zero mean, the expected mean error of the cost of *selected* projects will be larger than zero. In other words, the selected investments will exhibit systematic cost overruns, purely as a consequence of the selection process.

As this paper will show, all that is essentially required for selection bias to be present in project appraisal is that there is some kind of selection process in operation whereby selection is influenced by a noisy prediction. We describe the process in a stylised way in a model that comprises a noisy prediction step and a noisy decision step. This description fits easily with selection of projects from a list of projects. Such a description also fits with projects that seem to be of a more unique nature. Consider that projects that come to the attention of the public have generally been through a long, more or less formalised screening process. Initial estimates of costs and benefits may have been made at various stages in the selection process and numerous potential projects given up in the light of such information. Therefore the potential projects that come to our attention have already been selected; they are not random.

It is costly to appraise transport investment projects. For the largest projects, the traffic forecast alone may cost several hundred thousand Euros, and it takes a deliberate decision to incur such costs. Thus the fact that estimates of investment costs and benefits have been prepared implies that some kind of selection process will have been in operation. It is hard to imagine observing any list of projects that is not already heavily selected under influence of some preliminary prediction of costs and benefits.

Faced with systematic cost overruns and benefit shortfalls, a decision-maker may conclude that a stricter selection criterion is necessary. For example, several countries are implementing so-called “uplifts” in their procedures for project appraisal (as suggested in Flyvbjerg, 2008). However, we will show that raising the bar for project selection will *increase* the bias. If the same uplift is applied to all projects, it will obviously not affect project selection, holding the number of selected projects constant. Uplifts will only improve project selection if projects can be divided into classes, where the bias of each class is known in advance. However, uplifts may still be useful even if project selection is unaffected, since it enables more accurate estimates of the aggregate long-term budget consequences of decisions, if the size of the outcome bias is known from experience.

Given the large uncertainties in cost and benefit estimates that are observed ex post, it may seem tempting to draw the conclusion that cost-benefit analysis is useless as a tool for selecting which alternatives should be prioritized. In the words of Flyvbjerg (2009): *“With errors and biases of such magnitude in the forecasts that form a basis for cost-benefit analyses, such analyses will also, with a high degree of certainty, be strongly misleading. ‘Garbage in, garbage out’, as the saying goes.”* However, we demonstrate under quite general assumptions that the average selected project has a higher payoff than the average project, even if forecasts are uncertain. Hence, selection based on predicted payoffs is still beneficial in this sense, even with uncertain forecasts. Although the gain in average payoff from selection may vary, it is always positive regardless of how much noise there is in the predicted project payoff. In a numerical

illustration based on data on real-world transport investments, we find that the benefit-cost ratio turns out to be a robust selection criterion in the sense that the average benefit-cost ratio of the selected projects greatly outperforms random selection even for large uncertainties in benefit and cost estimates. Hence, the claim that “*cost-benefit analyses [...] will be strongly misleading*” is unfounded.

This paper begins in section 2 by formulating a stylised model of a project selection procedure. The model incorporates the essential elements of such a process, otherwise it imposes minimal structure. This ensures that conclusions will be applicable under a wide range of circumstances. In the model, projects are decided based on predicted payoff that is related to actual payoff. The actual payoff is observed only after projects have been selected and only for those projects that were selected. There is a selection mechanism that selects projects with a probability that increases as a function of the predicted payoff. The model assumes that *ex ante* predictions are unbiased. In spite of this, the model shows that the forecast error will be positive on average for those projects that were selected and a bias will seem to exist. Strengthening the selection criterion to compensate for the bias will actually increase the bias. The selected projects will have a higher average actual payoff than unselected projects.

Section 3 illustrates the empirical relevance of the selection bias using a database of projects. The list consists of the 461 road and rail investments of all types and sizes, that competed for inclusion in the Swedish transport investment plan 2010-2021. We show that plausible uncertainties in *ex ante* estimates of costs and benefits can give rise to large bias for selected projects. If the selection process is more competitive, i.e. more projects are rejected, average cost overruns and benefit shortfalls increase. Moreover, the examples indicate that the average benefit-cost ratio of the selected projects decreases only slowly as the noise in benefit and cost estimates increases. Hence, the benefit-cost ratio seems to be a useful selection criterion even when *ex-ante* estimates are highly uncertain.

2 A MATHEMATICAL MODEL OF FORECAST BIAS AS SELECTION BIAS

Projects are drawn from a population of projects. A project is characterised by a random variable X that represents the payoff of the project. It is not necessary to specify which payoff measure we are talking about. All that matters is that it is a measure that influences the decision whether to carry out the project.

The project payoff is not observed initially. It can only be observed if the project is carried out. At the time a project is decided, one observes instead the random variable Y that depends on X and represents a prediction of X . The relationship between predicted and actual payoff can be specified in a very general way by letting $Y = f(X, \varepsilon)$, where ε is a noise term that is independent of X and i.i.d. over projects, and where f is strictly increasing in ε .

The forecast error is $Y - X$ and we assume that forecasts are unbiased, which means that the forecast payoff is equal to the actual payoff on average: $E(Y - X) = 0$. This is weaker than assuming that $E(Y|X) = X$, which would mean that a regression of forecast payoff against actual payoff (if the latter could be observed) would yield the 45 degree line.

The decision whether to carry out a project is based on whether the predicted payoff Y exceeds a given threshold c , but it is not a deterministic process. The model allows for idiosyncratic factors in the decision process by saying that a project is selected if and only if $Z \equiv g(Y - c, \delta) \geq 0$, where δ is a random term that is i.i.d. over projects and independent of X and ε . The function g is assumed to be strictly increasing in its first argument, so increasing the value of c decreases the probability that a project is carried out. All random variables are assumed to have densities. For analytical simplicity, we assume that errors are supported on the whole real line and that f is such that for any number c and given X , $P(Y \geq c|X) > 0$.

The intuition behind our main results is simple, and is illustrated in Figure 1. Black dots are the true, unobserved payoffs of number of project suggestions – for the purpose of illustration, we assume that the true payoffs are the same for all projects. A decision maker observes forecast payoffs, illustrated by circles, and only projects above some selection threshold are realized. The true payoffs are revealed once projects have been realized. It is obvious that the true payoffs will on average be lower than forecasted for the *selected* projects, even if forecasts were unbiased for *all* projects. In other words, the selection process will mean that the decision maker will be disappointed by the true payoffs of the selected projects, on average – even if the forecast errors for *all* projects were zero on average; but the decision maker will never learn the true payoffs of the non-realized projects.




Figure 1. The selection bias mechanism. Black dots are true payoffs, circles are forecasted payoffs. Only projects above the selection threshold are realized.

It is also obvious that the more competitive the selection process is – the higher the selection threshold is – the larger will be the average difference between the estimated payoffs and the true ones.

This simple idea is formalized and made more general below. True payoffs do not have to be equal for the same phenomenon to occur, and the selection process does not have to be a strict payoff threshold: as soon as the selection probability is to some extent affected by the forecasted payoff, a bias will result.

2.1 Analysis

We now look only at projects that have been carried out. For these projects we observe both the predicted payoff Y and the actual payoff X . Thus, we observe draws from the

distribution of forecast errors conditional on selection. The probability that the forecast error is smaller than some number t , conditional on selection is denoted

$$P(Y - X \leq t | Z \geq 0).$$

Throughout we consider only situations where selection makes a difference, i.e. where $0 < P(Z \geq 0) < 1$. The following proposition states that the prediction error is positive on average for projects that are carried out. Hence it will seem as if predictions are biased even though they actually are unbiased ex ante. But the selection of projects is such that projects where the prediction error is positive are more likely to be carried out, everything else equal, and this effect produces the bias.

Proposition 1. *Payoffs are systematically overestimated for the realised projects: $E(Y - X | Z \geq 0) > 0$.*

All proofs are given in the Appendix. The next proposition states that the bias increases if the required payoff is increased. The intuition is the following. If the required payoff was very small, then its effect would be negligible and all projects would have almost the same probability of being carried out. In this case, there would be no bias. As the required payoff increases, the selection process becomes more important and it is that which causes bias.

Proposition 2. *A larger cutoff c implies a larger bias.*

It is of interest to seek testable implications of the model and it is particularly useful if such implications can be tested using a sample consisting only of projects that have been carried out. The next proposition states that the bias must decrease as a function of actual payoff, if the predicted payoff depends additively on actual payoff and noise.¹ This relationship can be tested in a regression of bias against actual payoff of selected projects.

Proposition 3. *If forecast error is additive $f(X, \varepsilon) = X + \varepsilon$, then the bias is decreasing as a function of real payoff: $\frac{\partial}{\partial x} E(Y - X | X = x, Z \geq 0) < 0$.*

When payoffs of selected projects are overestimated on average and raising the threshold for selection only makes the bias worse, it may be natural to ask whether selection actually does any good. The final proposition affirms, unsurprisingly, that selected projects do indeed have higher average payoffs than the average project, provided that the prediction Y is an increasing function of the actual payoff X . This is a natural requirement to impose; it is introduced at this stage only because there was no need for it before.

Proposition 4. *If $Y = f(X, \varepsilon)$ is strictly increasing as a function of X , then a selected project yields higher payoff on average than a random project: $E(X) \leq E(X | Z \geq 0)$. The gain from selection increases as the threshold is raised: $\frac{\partial}{\partial c} E(X | Z \geq 0) \geq 0$.*

3 SIMULATION RESULTS

¹ Given appropriate sign restrictions, additivity can be achieved from a multiplicative model using a log-transformation.

We have shown above that predicted payoffs will be biased as soon as projects are selected from an underlying pool of candidate projects, and that this bias will be larger the stricter the selection criterion is. How large the bias will be depends, however, on the noisiness of forecasts and the distribution of actual payoffs. We are not able to directly observe these quantities since we do not have available a database of forecasts and outcomes for both selected and unselected projects. What we can do to support our story is to use numerical examples to investigate the selection mechanism with plausible noise levels and an actual real-world distribution of payoffs. We will also explore the usefulness of CBA as a selection tool when cost and benefit estimates are uncertain, by investigating how the average benefit-cost ratio of the selected projects is affected by forecast uncertainty.

The numerical examples are based on all 461 suggested transport investments, shortlisted for possible inclusion in the Swedish Transport Investment Plan 2010-2021. The investments are described in Eliasson and Lundberg (2012). Each suggested investment has a total benefit B and an investment cost C and hence a benefit-cost ratio $BCR=B/C$. In the simulations, we will take these to be the true benefits and costs, unobserved by the analyst when selecting projects. The analyst only observes predicted benefits and costs B' and C' , where $B' = B * \varepsilon_B$ and $C' = C * \varepsilon_C$ and $\varepsilon_B, \varepsilon_C$ are random numbers generated in our simulation. The analyst is assumed to select the 100 projects with the highest estimated benefit-cost ratios², emulating project selection under a given budget constraint.³

We will follow the convention that the relative demand error is defined as $\frac{B'}{B} - 1$ while the relative cost error is defined as $\frac{C}{C'} - 1$. Note that the relative cost error is defined as the outcome divided by the prediction, while the relative benefit error is defined the other way around. This is consistent with much of the literature, and will ensure that estimates of benefits, costs and benefit-cost ratios are all unbiased. We assume that benefits are normally distributed while costs are lognormally distributed. With $\varepsilon_B \sim N(1, \sigma_B)$ and $\ln(\varepsilon_C) \sim N(\frac{\sigma_C^2}{2}, \sigma_C)$, we thus have $E\left(\frac{B'}{B} - 1\right) = 0$, $E\left(\frac{C}{C'} - 1\right) = 0$ and, in particular, that the forecast of the benefit-cost ratio is unbiased: $E(BCR') \equiv E\left(\frac{B'}{C'}\right) = \frac{B}{C} \equiv BCR$. Figure 2 illustrates true benefit-cost ratios BCR (black) plotted against predicted benefit-cost ratios BCR' (red), simulated for $\sigma_B = \sigma_C = 0.2$.

² This is of course a simplification of the real decision process. As stated in the previous section, it is not necessary that selection is decided on perceived payoff (such as costs and benefits) alone – it is enough that the perceived payoffs affect project selection to some extent, such that projects with higher perceived payoffs are more likely to be selected.

³ This selection process is equivalent to the one used in propositions 1 and 2, where all projects with payoff above a threshold were selected. This alternative selection process makes it easier to compare outcomes for different error variances, and also to vary the number of candidate projects.

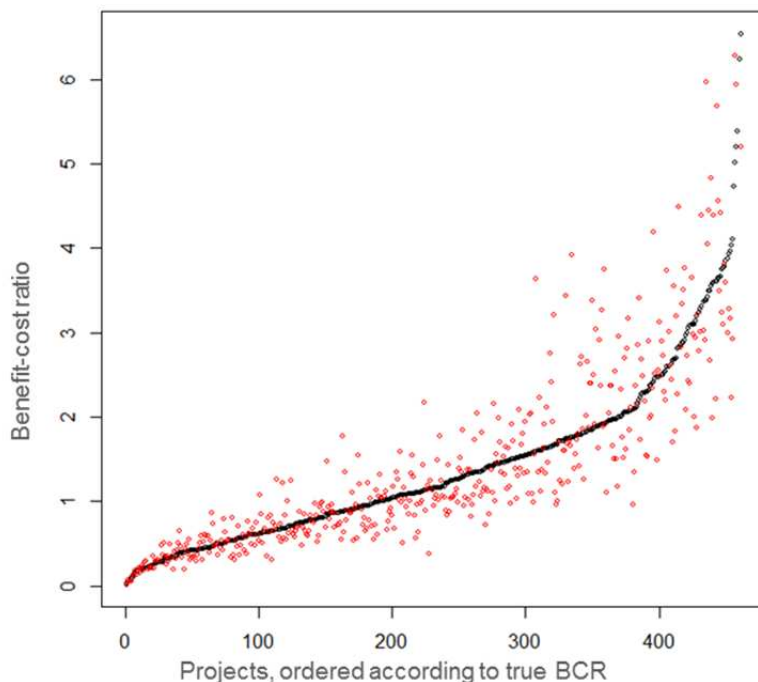


Figure 2. Simulated benefit-cost ratio estimates of the suggested investments (red) and true benefit-cost ratios (black), simulated for $\sigma_B = \sigma_C = 0.2$.

3.1 Cost overruns and benefit shortfalls

The first set of simulations illustrates how the selection increases with increasing noise in the predictions. For each simulation, σ_B and σ_C are fixed, and errors ϵ_B and ϵ_C are drawn for each candidate project. This gives B' , C' and BCR' for each project, and the 100 “best” projects (based on BCR') are selected. σ_B and σ_C are varied between 0 and 0.5 with a step of 0.01, repeating each step 20 times. Figure 2 shows how the relative cost and benefit errors of the selected alternatives grow as the standard deviation of the costs and benefits grows. Each dot represents one simulated selection, with the simulated standard deviation on the x-axis and the mean error of the selection on the y-axis. Note that the mean error of the selected projects is always positive, as predicted by proposition 1.




Figure 3. Mean relative cost error (left) and benefit error (right) of the 100 selected projects, plotted against the standard deviation of the estimated costs and benefits of all projects.

In this example, the relative errors are approximately proportional to the underlying standard deviations. An underlying standard deviation of 0.3 results in an average overestimation of costs and benefits of approximately 15%; if the underlying standard deviation is 0.5, the overestimation increases to around 30%.

The size of the resulting bias does not only depend on the uncertainties in the cost and benefit estimates; it also depends on how competitive the selection process is, and how much the true benefits and costs vary across projects. First, as shown in proposition 2, the bias increases the more competitive the selection process is: the smaller the share of selected projects is, the larger the resulting selection bias will be. This is illustrated in Figure 4.




Figure 4. Relative errors in costs (left) and benefits (right) of selected projects as a function of the fraction of projects that are selected.

In the simulations shown in Figure 4, the pool of candidate projects has been enlarged by adding it to itself (from 461 to 4610 alternatives), and a varying fraction is selected from this list of candidate projects. The resulting biases in benefits and costs are plotted on the y-axis against the fraction of projects that are selected on the x-axis. 20 simulation runs are made for each fraction, using $\sigma_B = \sigma_C = 0.4$. When 10% of the project suggestions are selected, the costs of the selected projects are underestimated by almost 40%, while benefits are overestimated by almost 30%. For lower selection fractions, the selection biases increase rapidly.

The second factor affecting the resulting bias is how much the true payoffs vary compared to the noise in the forecasts. Intuitively, if the best projects are much better than the average ones, the best projects will be selected even if forecasts are noisy, and hence the bias will be smaller. This is illustrated in Figure 5. Here, true benefits and costs B and C are replaced in the simulation by $B_\alpha = B^\alpha$ and $C_\alpha = C^\alpha$. The resulting bias in benefits and costs is plotted on the y-axis against α on the x-axis. The simulations are based on $\sigma_B = \sigma_C = 0.4$. As expected, the bias increases the more similar benefits and costs of underlying projects are, i.e. α tends to zero.




Figure 5. Mean cost (left) and benefit (right) errors of selected projects as a function of how similar benefits and costs of underlying projects are.

Lundberg et al. (2011) summarise the findings of 21 studies of cost overruns. Average cost overruns in these studies range from 0% to 50%. The simulations presented here indicate that plausible values of benefit and cost uncertainties, number of candidate projects and benefit and cost differences between candidate projects can easily give rise to bias of this magnitude. Obviously, this does not prove that intentional misrepresentation does not occur; but it shows that the observed magnitudes of cost overruns and benefit shortfalls do not prove that it does.

3.2 CBA as selection criterion when costs and benefits are uncertain

When uncertainties in cost and benefit estimates are large, using benefit-cost ratios as the selection criterion may seem dubious. Proposition 4 states that choosing the projects with the highest estimated BCR will still yield a better outcome on average than choosing projects at random. The size of the gain, however, is an empirical matter. In this section, we explore how the average BCR of the selected projects is affected by the uncertainties in the cost and benefit estimates.

Figure 5 shows the mean BCR of the selected investments (left), and the share of the 100 selected investments that in fact belong to the actual top 100 (right), i.e. the share of “correctly selected” projects.




Figure 6. Mean BCR of selected projects (left) and percent correctly selected projects (right) as a function of the relative BCR error. Red lines: mean BCR of all projects (left) and percent correctly selected projects (right) under random selection.

In this example, the mean BCR of the selected investments remains much higher than the average BCR of all investments, despite considerable noise in the predictions of costs and benefits. Moreover, the share of the actual best projects that are included in the selection remains high, despite the errors in the cost and benefit estimates.

4 CONCLUDING REMARKS

This paper has considered a process whereby selection of projects is influenced by some noisy but unbiased prediction of a payoff. Under very general circumstances, such a process will lead to selection bias, i.e. that the predicted payoff is smaller on average than the payoff observed ex post. The selection bias can easily be large at plausible noise levels.

It is important to realise that the results of this paper do not require that decisions are *determined* by predicted payoffs, only that noisy predictions have some influence on decisions. Still, it might be argued that cost overruns and benefit shortfalls are unlikely to be a result of selection bias because predictions of costs and benefits do not affect decisions. While there is evidence that benefits and costs do affect project selection (Eliasson & Lundberg, 2012; Nellthorp & Mackie, 2000; Odeck, 2010), other studies have found limited or no evidence of benefit-cost ratios affecting project selection (Nilsson, 1991; Odeck, 1996). However, it seems safe to say that predicted costs alone virtually always affect project selection, simply because resources are generally scarce. Moreover, as we noted in the introduction, there are in general very many possible projects that could be considered but never make it to the point where predictions of costs and benefits will be made and published. This early process is likely to be affected by some of the same factors that later might cause errors in predictions of costs and benefits. It therefore seems clear that there must always be a selection process that leads to the kind of bias that has been observed. Our simulations show that the selection bias on its own is enough to generate the magnitudes of bias encountered in reality, for plausible values of the relevant variables (see the survey in Lundberg et al., 2011).

Given the large uncertainty inherent in predictions of costs and benefits, one may question the usefulness of basing project selection on these. This is for example the conclusion in Flyvbjerg (2009). But as we have shown, this argument does not fly: if projects are selected based on predicted costs and benefits, even if the predictions are very uncertain, then the selected projects will turn out to perform better on average on these criteria than random projects. In fact, in our numerical investigations, the benefit-cost ratio is a surprisingly robust selection criterion even under considerable uncertainty, yielding much higher average benefit-cost ratios than random selection.

One of the suggested ways to remedy biased predictions is to use so-called uplifts, whereby predicted costs and benefits are corrected by the expected bias (Flyvbjerg, 2008). If the expected magnitude of the aggregate bias is known in advance, uplifts can be useful since they enable more precise aggregate budget planning. If it is possible to ascribe different uplifts to different classes of projects, using uplifts may also lead to better project selection. It should be noted, however, that imposing a stricter selection criterion, for example requiring a higher threshold benefit-cost ratio, will in fact increase the resulting bias in outcomes, contrary to the intention.

Of course, the demonstration that selection leads to selection bias does not rule out the existence of bias in the ex ante evaluation of investment projects. We have merely shown that it is not possible to conclude from the observation of ex post bias that there must have been bias in the predictions ex ante. We have also presented some numerical evidence that a selection process on its own is enough to generate bias of typical magnitudes, given plausible parameters. So while we can refute the argument of Flyvbjerg, we cannot refute his conclusion. Strategic misrepresentation by project promoters may well exist; but the existence of systematic cost overruns and benefit shortfalls does not prove this. As long as projects compete for selection based on uncertain, formal or informal, predictions of costs and benefits, these phenomena are bound to occur.

5 ACKNOWLEDGEMENTS

We are grateful to Ken Small for comments. Mogens Fosgerau has been supported by the Danish Strategic Research Council. Jonas Eliasson is supported by VINNOVA and the Transport Administration.

6 REFERENCES

- Benoît, J.-P., & Dubra, J. (2011). Apparent Overconfidence. *Econometrica*, 79(5), 1591–1625.
- Eliasson, J., & Lundberg, M. (2012). Do Cost-Benefit Analyses Influence Transport Investment Decisions? Experiences from the Swedish Transport Investment Plan 2010–21. *Transport Reviews*, 32(1), 29–48.
- Flyvbjerg, B. (2008). Curbing Optimism Bias and Strategic Misrepresentation in Planning: Reference Class Forecasting in Practice. *European Planning Studies*, 16(1), 3–21.

- Flyvbjerg, B. (2009). Survival of the unfittest: why the worst infrastructure gets built—and what we can do about it. *Oxford Review of Economic Policy*, 25(3), 344–367.
- Flyvbjerg, B., Holm, M. S., & Buhl, S. (2002). Underestimating Costs in Public Works Projects: Error or Lie? *Journal of the American Planning Association*, 68(3), 279–295.
- Flyvbjerg, B., Skamris Holm, M. K., & Buhl, S. L. (2004). What Causes Cost Overrun in Transport Infrastructure Projects? *Transport Reviews*, 24(1), 3–18.
- Flyvbjerg, B., Skamris Holm, M. K., & Buhl, S. L. (2005). How (In)accurate Are Demand Forecasts in Public Works Projects?: The Case of Transportation. *Journal of the American Planning Association*, 71(2), 131–146.
- Heckman, J. J. (1979). Sample Selection Bias as a Specification Error. *Econometrica*, 47(1), 153–162.
- Lundberg, M., Jenpanitsub, A., & Pyddoke, R. (2011). *Cost overruns in Swedish transport projects* (CTS Working Paper No. 2011:11). Centre for Transport Studies, KTH Royal Institute of Technology.
- Nellthorp, J., & Mackie, P. (2000). The UK Roads Review—a hedonic model of decision making. *Transport Policy*, 7(2), 127–138.
- Nilsson, J.-E. (1991). Investment Decisions in a Public Bureaucracy: A Case Study of Swedish Road Planning Practices. *Journal of Transport Economics and Policy*, 25(2), 163–175.
- Odeck, J. (1996). Ranking of regional road investment in Norway. *Transportation*, 23(2), 123–140.
- Odeck, J. (2010). What Determines Decision-Makers' Preferences for Road Investments? Evidence from the Norwegian Road Sector. *Transport Reviews*, 30(4), 473–494.
- Thaler, R. (1988). Anomalies: The Winner's Curse. *Journal of Economic Perspectives*, 2(1), 191–202.

van Wee, B. (2007). Large infrastructure projects: a review of the quality of demand forecasts and cost estimations. *Environment and Planning B: Planning and Design*, 34(4), 611 – 625.

Appendix

6.1 Proof of propositions 1 and 2

Consider the expected forecast error conditional on selection and use the law of iterated expectations to find that

$$\begin{aligned} E(Y - X|Z \geq 0) &= E[E(Y - X|Z \geq 0, X, \delta)|Z \geq 0] = \\ &= E[E(Y - X|Y \geq c + g^{-1}(0|\delta), X, \delta)|Z \geq 0] = \\ &= E[E(f(X, \varepsilon) - X|\varepsilon \geq f^{-1}(c + g^{-1}(0|\delta)|X), X, \delta)|Z \geq 0], \end{aligned}$$

where the second equality uses that g is invertible in Y for any δ and the third equality uses that f is invertible in ε for any X .

For any given value of X and δ , $E(Y - X|\varepsilon \geq f^{-1}(c + g^{-1}(0|\delta)|X), X, \delta)$ exists and is increasing as a function of c . Hence also the conditional expected forecast error $E(Y - X|Z \geq 0)$ is increasing as a function of c . This proves proposition 2.

Let now $c \rightarrow -\infty$. Then in the limit all projects are selected and $E(Y - X|Z \geq 0) \rightarrow E(Y - X) = 0$. This shows that $E(Y - X|Z \geq 0) > 0$ for any value of c as required.

6.2 Proof of proposition 3

To demonstrate the inequality, use that

$$E(Y|X = x, Z \geq 0, \delta) = x + E(\varepsilon|\varepsilon \geq c - x + g^{-1}(0|\delta)) = x + \frac{\int_{c-x+g^{-1}(0|\delta)} \varepsilon h(\varepsilon) d\varepsilon}{\int_{c-x+g^{-1}(0|\delta)} h(\varepsilon) d\varepsilon}.$$

Then differentiate with respect to x to find that

$$\begin{aligned} \frac{\partial}{\partial x} E(Y|X = x, Z \geq 0, \delta) &= 1 - \frac{\int_{c-x+g^{-1}(0|\delta)} \varepsilon h(\varepsilon) d\varepsilon}{\left(\int_{c-x+g^{-1}(0|\delta)} \varepsilon h(\varepsilon) d\varepsilon\right)^2} h(c-x) + \frac{(c-x)h(c-x)}{\int_{c-x+g^{-1}(0|\delta)} h(\varepsilon) d\varepsilon} \\ &= 1 - \frac{\int_{c-x+g^{-1}(0|\delta)} (\varepsilon - c(-x)) h(\varepsilon) d\varepsilon}{\int_{c-x+g^{-1}(0|\delta)} h(\varepsilon) d\varepsilon} \frac{h(c-x)}{\int_{c-x+g^{-1}(0|\delta)} h(\varepsilon) d\varepsilon}, \end{aligned}$$

which is strictly smaller than 1. Use last the law of iterated expectations to find that

$$\frac{\partial}{\partial x} E(Y|X = x, Z \geq 0) = E\left(\frac{\partial}{\partial x} E(Y|X = x, Z \geq 0, \delta)\right) < 1.$$

6.3 Proof of Proposition 4

Note that

$$E(X|Z \geq 0) = E(E(X|X \geq f^{-1}(c + g^{-1}(0|\delta)|\varepsilon), \delta, \varepsilon)) \geq E(X).$$

The second statement of the proposition follows from straightforward differentiation of

$$E(X|X \geq k) = \frac{\int_k^{\infty} xf(x)dx}{\int_k^{\infty} f(x)dx}$$

with $k = f^{-1}(c + g^{-1}(0|\delta)|\varepsilon)$ being an increasing function of c .