Notes on Estimation and Retrospective Estimation

Over the last few months I've been thinking about how people perform estimates on software teams. I've take time - not enough - to look through some of the research on estimation and try and understand how we can improve estimates.

Originally I was interested in "retrospective estimation" - how we quantify the time we spent on a task which has been already been completed. However, it seems the research on this topic is thin on the ground; my attention turned to "predictive estimation" - judging how long it will take to perform a task.

What follows are some very rough notes which result from reading several research papers, and some comments of my own. I intent to extract one or two (slightly) better written blog entries from this work and provide the full notes (the working drafts if you like) to a) better explain myself and b) in the hope that others will build on this work.

Note: I have bolded what I think are the key messages of this project.

Blogs and notes in the rough

These notes and observations are not intended, directly, for publication. Nor have they been through any form of review process. Please consider these notes *preliminary sketches*. Some of the information, even text, from this paper will appear in blogs and may well be incorporated into future journal articles.

Two blogs entries currently refer and summarise parts of this work:

- Humans can't estimate tasks, http://allankelly.blogspot.com/2011/03/humans-can-estimate-tasks.html

- Apology, correction and the Estimation Project, http://allankelly.blogspot.com/2011/03/apology-correction-and-estimation.html

Some of the research findings will also be incorporated into courses offered by the authors company, Software Strategy. Please see the Software Strategy website for details of course - http://www.softwarestrategy.co.uk/training.html.

A correction and an apology

In January (2011) Jon Jagger sat in on one of my Agile training courses. Afterwards he dropped me an e-mail to get some more information on a few points and I was glad to answer him.

Jon asked me about two references I appeared to cite. Now if someone, me or anyone else, goes around saying "there is evidence for this" or "this is backed by research' they should be able to back it up. I have failed on this count. I apologize to Jon and everyone else who has heard me say these things. I can't track down my references. I should not have claimed there is research which I can't provide a reference for.

Trying to track down these references led me to an interesting little research project I've been calling "The Estimation Project". More of this project and my findings later, for now I want to look examine my claims.

I'm going to repeat my claims now for two reasons. First, I'm hoping that someone out there will be able to point me at the information. Second, if you've heard me say either of these things and claim there is research to back it up then, well I apologise to you too.

The first thing Jon pulled me up on was in the context of planning poker. I normally introduce this as part of one of my course exercises. Jon heard me say "Make your estimates fast, there is evidence that estimates fast estimates are just as accurate as slow ones."

Does anyone have any references for this?

I still believe this statement, however, I can't find any references to it. I thought I had read it in several places but I can't track them down now. I'm sorry.

Thinking about it now I think I originally heard another Agile trainer say this during a game a planning poker a few years ago. So its possible that others have heard this statement and someone has the evidence.

Still, I think there is value in playing planning poker fast because it allows you estimate a lot of work rapidly. Plus, having a whole team estimate work is expensive, so to minimise the cost I'd like to see it done fast.

Now at this point someone will say "Why not have an expert make the estimate?" Well there is research that shows that **several independent estimates will give a more accurate estimate than one individual**, irrespective of how much experience any one person has. i.e. You are better off with several people than one expert. (This is basically the Wisdom of Crowds argument; for something shorter see "Why Forecasts Fail" in the MIT Sloan Management Review (Makridakis et al., 2010).)

Turning to my second dubious claim, I said: "I once saw some research that suggests we are up to 140% in judging actuals." It is true that I read some research that appeared to show this - I think it was in 2003. However, I can't find that research so I can't validate my memory, therefore this is not really valid.

By "actuals" I mean: the actual amount of time it took to do a piece of work.

Again, does anyone know the research I have lost the reference to?

And again, I apologise to Jon and everyone who has heard me state this as fact. It may well be true but I can't prove it and should not have cited it.

Why did I say it? Well, I find a lot of people get really concerned about the difference between estimates and actual, whether some task was completed in the time estimated. To m mind this is a red herring, actuals to my mind only get in the way.

The way I teach teams to estimate and track work is based on the XP velocity idea. At the end of each iteration teams count how many units of work that

have done and graph it. They then base the amount of work they plan to do on the average of the last few iterations. Thus the system is self-adjusting.

I always tell classes to always work with estimates and don't bother with actuals, i.e. a task is only estimated in advance and no track is taken of how long it actually takes.

This is where it starts to get interesting and The Estimation Project starts. In trying to track down this research I did find some other interesting research which I will now describe and include in future blogs.

The first thing to note is what I have been called "actuals" are actually estimates which are performed retrospectively. The little research on this goes by the title: "retrospective estimation."

An short, amateur, literature review

Scientists often start their research with a literature review, indeed, some scientific papers are only literature review. In order to get to the bottom of the estimation question, or rather the retrospective-estimation question, I decided to conduct my own little literature review.

I've endeavoured to read a lot of research about estimation - both forward and retrospective estimation/prediction. I sourced the papers I could from the internet and from the Ebsco repository. However I know there are more papers in other repositories which I cannot access.

"Cannot" is a subjective term, in use it knowing I could have found a way to access more papers. I could beg membership at a local University library, I could visit the library at the University where I did my Masters degree, I could camp out at the British library, I could buy subscriptions and papers from publishers but these would all be undertakings in terms of time or money.

One insight from this process: what the academic community regards as common practice isn't practical for practioners. These papers are behind firewalls, sometimes difficult to find, they cost to access, can be hard to read and don't always agree with each other.

Its very each for someone - academic, trainer or practioners - to say "Use fact based methods" but determining what the facts are isn't easy.

So maybe I was lazy, although I've spent several days (cumulatively) reading papers, I could have spent more. I was happy to spend a little money but scientific papers are not cheap, besides, between the Internet and Ebsco I could get more than enough papers to look at.

Next I should say my literature review skills are somewhat limited. And they aren't even as good as when I did my Masters degree. So, from an academic stand point you could question my research. Had I been conducting a through literature review I would have reviewed all the citations in these papers and read a good few of those papers too.

I will admit, I didn't read every word of every paper. I read lots of each paper - and some I did read entirely - but I concentrated on top-and-tail, that

is opening summary, conclusion/discussion and I missed out most of the hard research description in the middle.

Finally, I undoubtedly suffered from confirmation bias. There were things I was looking for and I probably overlooked things I didn't want to find. That said, I've tried to have an open mind.

Despite these caveats I think my research is significantly more than most bloggers do before they blog. And I'll admit, on this occasion it is a lot more than I normally do myself.

What follows here are my, rough, working notes. I publish them for completeness, so someone interested in this topic can find out more, and in the hope that someone will further this investigation.

Planning Fallacy: Kahneman and Tversky

One of the key pieces of work in this field is "Intuitive Prediction: Biases and Corrective Procedures" (Kahneman and Tversky, 1979). While a little old most of the other works in this field build on this research, support and extend the findings. The same authors are also the originators of "Prospect Thoery" which analyses decisions under risk. Prospect Theory looks interesting but I haven't had time to read about it yet.

(Interestingly Kahneman a psychologist an winner of the Nobel prize for Economics.)

This paper introduced the term **"Planning Fallacy"** to describe the inability of forward looking plans to accurately estimate the time work will take.

Perhaps unsurprisingly the authors find that **people tend to underestimate the amount of time that work will take to get done**. This isn't occasional or random, **its systematic**. The problem occurs again and again. The authors give an explanation of why this might be so.

Another finding in this paper is that **people are overconfident in their own predictions**. Even when confronted with their own mistakes people maintain their future confidence. People remain confident even when told that the majority of predictions are wrong.

The authors recommend that in order to get more accurate predications that experts work with analysts to explain their reasoning and engage in a dialogue.

Kahneman and Tversky talk specifically about "experts" doing the forecast but there is other, later (e.g. "Why Forecasts fail. What to do instead" and The Wisdom of Crowds), research that suggests that experts are no more likely to be accurate than anyone else. True they may have a better understanding of what needs doing but that doesn't mean they will be more accurate.

There is also an interesting aside that might be worth further investigation: people are actually good at predicting the frequency of reoccurring events. For example, weather forecasters are quite good at predicting tomorrows temperature because this prediction is largely based on frequency.

Hofstadter's Law

If the planning fallacy sounds familiar then it might be because it is another statement of Hofstadter's Law:

> "It always takes longer than you expect, even when you take into account Hofstadter's Law."

Coined by Douglas Hofstadter in his book, *Gödel, Escher, Bach: An Eternal Golden Braid (Hofstadter, 1980)* this is a less scientific analysis of the same thing. Intuitively Hofstadter's law makes sense, even sounds familiar, but it is the academically sound Planning Fallacy that has proof.

Francis-Smythe and Grove

Continuing my research on estimation I'd like to turn my attention to "On the Relationship Between Time Management and Time Estimation" (Francis-Smythe and Robertson, 1999).

This paper is interesting because it looks at time management not just estimation. One of they key points the authors make is: **time management is not just about estimation, it requires one to both plan and schedule and keep to a schedule.** Thus, time management required a) forward looking estimation, b) time monitoring when doing the task and c) review of time spent (retrospective estimation).

People who believe they are "good time managers" are actually better at estimating how long future tasks will take. This might be because once the estimate is made these people manage their time to accomplish the task within their allotted time. However, these same people are not very good at estimating time as it is passing.

The research suggest that the amount of control the estimator has over the task in hand plays a role in the accuracy of the estimate. Those who are in control of their own time are likely to give more accurate estimated - possibly because those making the prediction also choose the strategy for approaching the task.

Another factor in the accuracy of estimates was about routine. People with more routine seem to be better at estimating then those who aren't. I would conjecture that this finding fits with ability to forecast frequency mentioned by Kahnemann and Tversky. When events occur in a routine way they take on aspects of frequency.

The research also note that the evidence on people being able to estimate accurately is somewhat inconclusive. This isn't helped by the fact that some studies look at very short time periods (seconds and minutes) while other research looks at long time periods (hours and days).

These researchers find most people underestimate the time a future task will take a few people grossly overestimate the time the task will take.

This research also has a little to say about retrospective estimation. It seems that determining how long was spent on a task is just estimation. There

doesn't seem to be much evidence in this paper that retrospective estimation isn't a lot more accurate than forward estimation.

Zackay and Block

Next in my review of research on estimation is "Prospective and retrospective duration judgments: an executive-control perspective" (Zackay and Block, 2004). This paper discussed the subject I was originally interested in, retrospective estimation, but it does not shed much light on the subject.

These researchers are concerned with the mental processes which underlie our ability to retrospectively estimate how long a task took. What they find is that our brains make this judgement differently when we know in advance that we must report duration (i.e. we are told before hand that when the task is done we will be asked "how long did it take?") and when we don't know (i.e. after the task someone asks "how long did it take?" without warning.)

There is actually little in this paper about the correlation between reported time and actual time spent. The paper does say that **reported time was always less than the actual time**, so it would appear the planning fallacy holds for retrospective estimation.

Things improve a little when we know in advance that we must report the time spent, that is, the reported is longer than when we don't know in advance. But the reported time spent is still less than the actual time spent, i.e. the reported time is closer to the actual time.

However, the paper does conclude that when we know in advance that we must estimate the time spent we use some brain power to monitor time. When the task in hand consumes a lot of brain power then time reported converges with the "not knowing in advance" scenario. i.e. knowing in advance that we must report the time spent has less significance when the task is complex.

One more factor seems significant: the time gap between completing the task and reporting the time. The longer the gap the less difference knowing in advance makes.

Put this into a (software) work situation: does knowing that we must report the time spent on a task make a difference to the time we report?

The answer would seem to be Yes - although which is the most accurate isn't completely clear. But, if the task in question requires a lot of thinking it makes less difference.

Since most software development tasks are complex, actually knowing in advance that we must report our time doesn't make a big difference.

Also, I would suggest that while developers might know they have to complete a timesheet at the end of the week the fact that they need to report time is frequently forgotten about until it is asked for.

Conclusion here: **if you want to track actual time spent then always remember you need to report the time and record it as soon as the task is done**.

Cardaci's Mental Clock

Next up is "A Study of Temporal Estimation From the Perspective of the Mental Clock Model" Carmeci, Misuraca and Cardaci. (Carmeci et al., 2009).

Cardaci, has proposed a "mental clock model" which this paper investigations. Reading the discussion section of this paper it is clear that **researchers don't really understand how the mind works when retrospectively judging time**. The authors point out that there are several models and several studies which all suggest different ideas.

Cardaci's model is counter to what I think most people would believe intuitively: *if a task need more attention and mental work then time appears to slow down, while when tasks require less attention and mental processing time appears to go faster.* Perhaps not surprisingly the research in this paper supports this model.

Now, anyone who is reading all these papers, or even just my summaries, will notice something here. This finding contradicts the Zakay and Black finds which suggest that more mentally intensive tasks will slow down the perception of time.

On the face of it this doesn't move me any further forward to understanding whether retrospective time estimates - and tracking - is accurate or not. The "experts" don't have a conclusive answer.

However, the fact that even experts in **scientific experiments can't produce consistent findings**, or come up with a model to explain our reasoning process, does not bode well for corporate time logging systems which are administered by amateurs in decidedly unscientific conditions.

Buehler, et al.

Next up in my mini-literature review of estimation research is "Exploring the 'Planning Fallacy': Why People Underestimate Their Task Completion Times" (Buehler et al., 1994). This was the most detailed and statistics driven piece of research that I look at. The researches conducted five different experiments in all. Of all the papers I looked at this is probably the most insightful, although its also one of the most difficult to read because of all the statistics.

The abstract to this paper seems to say it all. The researchers test three hypothesis which they show to be true:

a) People underestimate their own but not others completion times

b) People focus on plan-based scenarios rather than relevant past experience when predicting

c) People undervalue past experience

There is an interesting aside in the opening pages of this paper. Some 15 years elapsed between the original planning fallacy paper and this paper (1994) yet the researches comment that the subsequent research into estimation was surprisingly sparse. This comment might go some way to

confirming what I myself found: the research in this area is not of the volume or certainty I expected.

Some findings from the studies.

- When asked to give Best-case and Worst-case estimates both estimates were equally wrong.   Pessimistic estimates were likely to be longer than optimistic ones but still underestimated the time required.

- There is a correlation between predicated time and actual time indicating under estimation is systematic not random. So it would seem people can predicate the relative magnitude for a task but that the actual magnitude estimation is too small. This would also imply that people can tell which tasks will take longer than others.

- There was no evidence that when people considered their own and other peoples past experiences estimation accuracy improved, however, this was in part because so few people referenced the past.

- When study participants had an external deadline they met the deadline in over 80% of cases - although their estimates were still wrong they and tasks too longer than expected they were done by the deadline.

- In all the studies participants exhibited high levels of confidence in their estimates.

The authors note that in one casual survey 73% of students admitted to finishing work the same day as a deadline. It seems that deadlines might exert the same, or even more, power than estimates. This would seem to support the Francis-Smythe and Grove suggestion that time-management also plays an important role.

The authors here go on and explore this a little further. In appears that when deadlines are present people change their behaviour in estimating. And it seems **deadlines are more significant than estimates in determining when work will be complete.** In one study 43% of participants completed a piece of work within their own estimate (not bad actually, I'm surprised how high it is) but 75% of participants completed it by the deadline.

In one experiment the research set two groups the same task but with different deadlines. The ones with the later deadline provided larger estimates. Yet the extra time did not make a difference to the actual time taken to complete the task.

Interestingly, although deadlines brought work to a completion on time few people admitted to considering deadlines when estimating their work. However, time estimates were highly correlated to deadline so it would appear that deadlines do indeed influence estimates even if we don't admit it.

One other twist these research came up with concerned past experience. If subjects were asked to "consider past experience" when making estimates 38% of subjects completed work within the estimated time, up from 29%. However, if subjects were asked in detail about past experience and helped to relate the new task to previous experience then 60% of subjects finished the work within the estimated time.

Curiously, although these figures make it look like recalling past events makes for better estimates they don't. The correlation between time spent and actual time was not significantly better for any of these three groups. This seems a little paradoxical: recalling previous events mean the estimated time is more likely to long enough to do the task, but it doesn't make the estimate any more accurate. Forcing people to **recall past events removed the bias towards optimism but it substituted pessimism.**

From this I would suggest that people who recall past events are more likely to "pad" the estimate, i.e. add some contingency just in case. So following this tactic may get lead to a "worst case" estimate time.

As an extra experiment these researchers also look at how third-party, an observer, estimates related, i.e. if someone estimates how long it will take me to perform a task.

Third-party observers tended to produce more pessimistic estimates, they also tended to use more of the available information when estimating and considered potential future problems more often.

Where the estimates better because they were done by a third-party? Sometimes, **in some scenarios the third-party estimates produced more accurate (closer to actual) estimates and in other scenarios the estimates were less accurate**. This does conflict a little with one of the other studies that suggested that third-party estimates were more accurate as a whole.

In conclusion this study finds:

• People underestimate how long it will take to perform a task

• People make estimates by mentally planning; not by considering the past

• There is no compelling evidence that using past performance will improve estimate accuracy

• While optimistic bias can be reduced accuracy cannot. (Or at least, these researches haven't found a way to do so yet.)

The researches also offer one interesting suggestion as to why humans are biased toward optimistic bias: as long as the deadline it met, the cost of being underestimating (being short) is minimal, indeed the motivational benefits of being optimistic are worth having. Pessimistic estimates on the other hand are likely to be demotivating. Even if tasks are not completed by the optimistic estimate they may be completed sooner than they would have been if an accurate estimate, using all available information, had been given.

Conclusion

So, from all this research I draw the following conclusion:

1. Prediction, and specifically estimation, isn't a particularly well-understood field. The researchers are still experimenting and producing models for how human brains cope with these tasks.

2. Human's aren't very good at estimating how long it will take to do something.

3. Humans are hard wired to underestimate how long it will take to do work. You can drive this out but you replace it with overestimates not accuracy.

4. Humans can estimate tasks relative to one another: underestimation is to be systematic; there is a good correlation between estimated time and actual time. In general humans can tell which is bigger and which smaller, i.e. relative magnitude of tasks.

5. Retrospective estimation isn't very accurate either; the "Planning Fallacy" seems to hold retrospectively. You might be able to improve it by putting more real-time effort into recording it but you will also increase the cost of recording it

6. Deadlines are more significant in when determining work will be done than many of us realize, or would like to admit.

There is another conclusion that is not completely proven but fits with the "groups estimate better" argument. Some people over estimate, and when people consider the past in detail they overestimate. In a group setting the few who overestimates may compensate a little for the mass who underestimate.

Some conclusions from the conclusions:

• If we ask people to estimate how long a piece of work will take they will probably underestimate. So if this estimate is then used as a deadline the deadline may well be missed.

• If people are encouraged, coerced, or scared into giving pessimistic estimates the estimate will be too long. If this estimate is then used as a deadline work is likely to be completed inside the time but there will be "slack" time. The actual time spent on the task may be the same either way but the total elapsed (end-to-end) time will be longer.

• Either way turning time estimates into deadlines doesn't seem like a good idea. But, deadlines themselves do seem like a good way of motivating people to complete work.

• There seems no reason to believe retrospective estimation is significantly more accurate than future estimation.

One of the key findings I am taking away from this reading is that the correlation between estimated time and actual time shows underestimation is systematic. **Therefore any time estimate is likely to be too small but in relation to other estimates the estimate is good.** You could write this as an approximate equation:

$$\text{Actual Time} = K \times \text{Estimated Time}$$

However, we I don't know, what I can't tell from the studies, is how constant K is. In my own mind am certain it will differ between individuals and between teams, thus it will differ between projects. I would also expect it to change over time.

So the best I think we can say is: K is approximately constant for a given individual or team in a given context over the short run.

One other thing which isn't considered in any of the studies I looked at, is: how long did it take to make the estimates? I would guess that someone who is being asked to think about, and record, past experiences, when making an estimate is going to take longer to make the estimate. So the overall time and cost of estimating is also going to increase.

And Agile....

As to practical implications for software development (and Agile)....

Broadly speaking I find this supports the way I teach teams to do work breakdown, estimation and deciding what to put in an iteration. Specifically:

- Work is broken down by the team who will do it: they have control over their own

- Work breakdown and estimation is partly (even largely) a design activity because the approach to the work is being discussed

- Estimation is done by several people, experts and non-experts alike using abstract points

- The amount of work which is scheduled for an iteration is an average of the last few iterations: this means the iteration capacity floats and is self adjusting

- Agile routines, iteration, the regular schedule of meetings and events, helps planning.

See Two Ways to Fill an Iteration for more details

I can't be sure about this but in estimating work I think there is an element of frequency estimation, the problem has been changed from an estimation of size to one of frequency. I'm not qualified to be sure about this but I can see the potential.

Turning to "actuals". I tell teams to ignore actuals and work with estimates. I believe these are like apples and pears so you shouldn't compare them. Additionally some teams and individuals spend a lot of time estimating time spent and even arguing over it. I don't see any value in this.

When planning an iteration I always work with estimates. Because the capacity level is allowed to float actuals aren't important.

The papers I have reviewed do not lead me to change my position. "Actuals" are just another estimate, albeit a retrospective one. I may not have found my "140%" figure in these papers but I see little reason to believe retrospective estimates are accurate enough to add anything significant to the discussion.

The Buehler, Griffin and Ross paper shows that people could estimate relative magnitude of tasks, and there was a high correlation between estimates and actual time. Thus, given that both estimates and actuals are wrong but correlated then why not just us the estimates and save the effort of mucking about with actuals.

Timesheets

Now I know some people are asked to fill in time sheets by their employers, and some of these time sheets are used for billing purposes or for tax credit claims. There are lots and lots of websites advertising time tracking systems and claiming accuracy. But they all - OK the vast majority - seem to rely on people entering time themselves. However, there is surprisingly little research on how accurate the data entered into these systems is.

If you need to engage in this fantasy then so be it. I sympathise but I don't see any point in fighting the system, my advice:

- For greatest accuracy fill these timesheets in as soon as you've done the work.

- In general fill in these sheets rapidly. Don't waste your time on a pointless exercise.

And to managers, accountants and financial controller who think my advice is flippant I say: *Go and research the subject yourself, your current methods don't work; they are an illusion.*

If you must have accurate numbers than go and find a real solution. However, most of you don't need accurate numbers, you only need to perpetuate the illusion up the chain and to the taxman.

Future research

As with all the best scientific studies I want to conclude with some recommendations for further research. It would be nice to imagine that a professional researcher will decide to tackle these questions but I imagine that it will depend on me finding more time.

- **Cost of estimation and accuracy**: none of the papers considered looked at how much time (and thus money in business) it cost to make estimates. It would be interesting if any correlation exists between more time spent making an estimate and ultimate accuracy. The next question to ask would be: is it justified to spend more time making estimates?

- **How constant is the systematic underestimation?** Is it as simple as: Actual Time = K x Estimated Time? How constant is K? Under what circumstances does it change?

- **Does the planning fallacy really hold for retrospective estimation? Are retrospective estimates in any way more accurate?** It looks like the planning fallacy does hold and there is little evidence that retrospective estimates are more accurate. However these are questions that need more research.

- **Retrospective estimation in the workplace**: Outside of the scientific environment, in the messy world of work - full of disruptions, distractions, time delays and everything else - how accurate is retrospective estimation? I suggest not very.

- **Accuracy of time reporting systems**: Building on the previous suggestions someone should really look at time reporting systems. Is there any reason to believe they are, or even could be, accurate?

That final question is likely to prove controversial. A lot of companies are making money from selling time reporting systems. They have a vested interest in research that shows their systems are accurate. So maybe they would find research. But, I don't expect any of them will fund research that might show the opposite. The risk is too high.

My belief is: humans are as bad at retrospective estimation as they are at predictive estimation, therefore any serious study of the accuracy of time reporting systems would show them to be flawed. Therefore, the companies that produce the systems will not invest in research.

In total speculation, the lack of such research (funded by time reporting companies) might even be taken to demonstrate that these systems are flawed.

In this context it is interesting to note that Capers Jones suggest that corporate time/money tracking systems are typically 70% inaccurate and are a major contributor to poor scheduling in companies (Jones, 2008). of course, before you can believe Jones suggestion one has to ask, how does he know the real figures in order to say system are inaccurate.

Finally, one more research point, actually, it is partly about research and partly about interpretation. There is a serious need for someone to look at the research, and conduct more specific research, in the software/IT development context.

None of the research I looked at was specific to the software development environment. Much of it can, and will, apply, but within the specific context it needs detailed examination. In addition, more specific research needs to be undertake to see how software estimation really stacks up.

Unfortunately, this requires more research skills, and time, than I have available. So I have a little wish, a hope.

It would be nice is to see some work on this subject specifically in the software field. If anyone is looking for a PhD research topic can I suggest: Obtaining useful estimates for software projects?

Personal insights

I think its worth recoding some personal remarks on this little "Estimation Project" of my own. When Jon Jagger first challenged me to show the evidence I thought "easy, I've got it on the hard disk". Once I realised I didn't I thought "arh well I'll just find in on the net" and I didn't.

But this did set me looking into retrospective estimation more and I decided to devote some time to it. Retrospecitve estimation research became estimation research and the whole thing has ended up taking a lot longer than I expected. The planning fallacy.

In fact I'm only pulling the whole thing together now because I'm fed up of the project sitting on my to do list half finished.

My second observation has to do with what I found. I am a little worried that little of what I found contradicts what I thought I knew, i.e. what I believed. There are several possible explanations for this. The first one is very self-congratulatory: I was right all along. I might have lost my reference but I was broadly right.

The second explanation is worrying and potentially casts a doubt over all my findings. Confirmation bias: I've only paid attention to the research that supports what I already believe. One way or another I've not found contradictory research.

There is one thing I have learned, one thing I will do differently as a result of this investigation.

In future when I ask people to estimate a piece of work, whether in a classroom exercise or in the wild I will request: "Please make an estimate based on how long you think it will take someone else to do the work."

This might improve the estimate accuracy; it will most likely make them slightly larger (i.e. pessimistic); it will certainly side-step the "I can do this in an hour but I think others will need a day to do it" scenario.

Finally, while I understanding a bit more about retrospective estimation I still haven't got a conclusive answer.

References

BUEHLER, R., GRIFFIN, D. & ROSS, M. 1994. Buehler, Griffin and Ross *Journal of Personalty and Social Psychology,* 67**,** 366-381.

CARMECI, F., MISURACA, R. & CARDACI, M. 2009. A Study of Temporal Estimation From the Perspective of the Mental Clock Model. *The Journal of General Psychology,* 136**,** 117-128.

FRANCIS-SMYTHE, J. A. & ROBERTSON, I. T. 1999. On the relationship between time management and time estimation. *British Journal of Psychology,* 90**,** 333-347.

HOFSTADTER, D. R. 1980. *Godel Escher Bach : An enternal golden braid,* Harmondsworth, Penguin Books.

JONES, C. 2008. *Applied Software Measurement*, McGraw Hill.

KAHNEMAN & TVERSKY 1979. Intuitive Prediction: Biases and Corrective Procedures. *TIMS Studies in Management Science***,** 313-327.

MAKRIDAKIS, S., HOGARTH, R. M. & GABA, A. 2010. Why Forecasts Fail. What to Do Instead. *MIT Sloan Management Review,* 51.

ZACKAY, D. & BLOCK, R. A. 2004. Prospective and retrospective duration judgments: an executive-control perspective. *Acta Neurobiol Ex***,** 319-328.