Neutral hours: a tool for valuing time and energy

Owen Cotton-Barratt*

We get lots of opportunities to convert between time and money, and it's hard to know which ones to take, since they use up other mental resources. I introduce the *neutral hour* as a tool for thinking about how to make these comparisons. A neutral hour is an hour spent where your mental energy is the same level at the start and the end. I work through some examples of how to use this tool, look at implications for some common scenarios, and explore the theory behind them.

1. Introduction

This is an introduction to my thoughts on the personal prioritisation problem of how to value time. I present a model which is simple enough to use for some personal decisions. It's also relevant if you need to value other people's time, for example if you want to do cost-benefit analysis for some intervention which will save or cost people time. In this section I give an introduction to the method; in subsequent sections I give more detailed thoughts.

As we have to make decisions all the time about how to spend time or money, it's useful to have an idea of the appropriate conversion rate. Indeed it's reasonably common advice to know what value to put on your time.

If you take the naïve approach of picking a value for your time and applying this to make decisions, this predictably gets things wrong in some circumstances, because of confounding factors that mean that not all your hours are equal; some are much more tiring than others. But simply saying "remember to account for confounding factors" isn't terribly helpful as practical advice if it doesn't tell you *how* to account for them.

My answer is a first-pass approach at getting these corrective factors right. It leverages our intuitions about picking trade-offs between time and energy to give a rule that is simple enough to think about and apply on-the-fly. I don't imagine it will let you always reach the optimal answer, but I think it comes closer: it will much more often get the correct answer when one course of action is substantially better than the alternative.

The central idea is not to count hours spent on an activity by the clock, but to weigh them according to how draining or recuperative they are:

* St Hugh's College, Oxford & Centre for Effective Altruism. Email: cotton-b@maths.ox.ac.uk

> A *Neutral hour* (Nh) is an hour in which you feel as tired at the end as at the start.

Now we can say that the *Neutral time cost* (NTC) of an activity is the number of neutral hours you must spend to complete it. If the activity is net-neutral for energy levels, this will be equal to the actual time taken. For net tiring activities it will be higher, and for net recuperative activities it will be lower.

1.1 Advice

Choose a value of your time, in $/Nh, and use this to make decisions. The idea is that it provides a guideline threshold: if an opportunity has a notably better exchange rate it's worth taking, and if it's notably worse, it's not. When the exchange rate you work out is very close to your threshold, it may be worth having a second look at your estimates, as more precise figures might change the decision.

This should prevent you from making dominated decisions, and you can look at tweaking the value used up or down to something more appropriate after you have tried it.

If the main reason you work your job is for the pay, you may want to try a value based on your hourly earnings. If so, it's important to remember to:

- Take the earnings for the marginal hour, if this is possible;
- Remove tax, so you are reckoning in take-home pay;
- Work out the NTC for the marginal hour of work. If an extra hour of work needs an extra hour of leisure the NTC is 2 hours.

These factors can all reduce the $/Nh of work, perhaps giving a figure which is significantly lower than the apparent value of an hour's time.

Something that pushes the other way is that for many people there are significant benefits to the job beyond take-home pay, which can push up the value of time. This can be true even if you are in a career primarily for the earning potential. For example a major benefit of working for a year in investment banking is the opportunity to work subsequent years in investment banking. This benefit could be noticeably larger than the value of even a high salary!

1.2 Simple worked examples

Alice works as a doctor. In addition to her regular contractual work, she can pick up extra hours at a higher hourly rate of $100/h, but which do not help her career. She could take more of these hours than she does, so she decides to base the value of her time on this for now.

Her take-home pay (after tax) is $70, and she finds the work somewhat tiring, needing an hour to recover per hour of extra work. So she takes as the value of her time $70/2Nh = $35/Nh.

She can use this figure to decide whether it's better to take up other casual work with different rates of pay and energy demands. She can also use it to evaluate time-saving opportunities.

For example she has a choice between taking:

a. A 3 hour bus journey costing $20;

b. A 2 hour train journey costing $60;

c. A 90 minute taxi journey costing $70.

She does not think she can work effectively on the bus, but the journey will be as restful as an hour's leisure, so the total NTC is 2Nh. On the train she thinks she can work on some reading, achieving what would normally take her 1.5Nh, so the total NTC is 0.5Nh. The taxi will also not allow work, but is more restful per minute than the bus; an hour's worth in total. So the NTC is 0.5Nh.

The train and the taxi actually have the same NTC, so the cheaper option (the train) is obviously preferable. And the train costs $40 more than the bus but saves 1.5Nh, which is cheaper rate than her threshold of $35/Nh, so she decides to take the train.

Note that if she'd taken the times and prices at face value, she'd reach a different conclusion. So was the framework of neutral hours necessary to get to the right answer? No, you could also get there by considering appropriate counterfactuals like:

a. Take the train, spend 45 minutes working overtime and 1 hour 45 minutes resting (4.5 hours, $60 – (0.75*$70) = $7.5 cost, net 1 hour recuperation)

b. Take the bus, and spend 1.5 hours doing the reading and resting up. (4.5 hours, $20 cost, net 1 hour recuperation)

Assuming you will be doing the reading, and taking some rest at some point, this gives a dominance argument for the train in favour of the bus. A similar case could be constructed to show that the train is better than the taxi. However, it's a slightly awkward calculation and certainly not something that is practical to perform as part of a quick calculation. Having the concept of neutral hours at hand made it much easier to arrive at the correct answer.

This is the strength of the method: it is not an ultimate theory of the value of time, but a pragmatic rule of thumb for improving personal decision making.

2. Why to value your time

2.1 Dominance arguments – potential for strict improvements

We often have opportunities to exchange time for money, or vice-versa. These can come at a variety of exchange rates. We take some of the opportunities but not all of them. If you pass up an opportunity to spend an hour to gain $X, and also pass up an opportunity to spend $Y to gain an hour, where Y < X, then your course of action is, at least in a simple model, dominated by the one where you take both opportunities and end up with the same amount of time but $(X-Y) more.

Still in the simple model where activities only have effects on time and money, if you have chosen a course of action which is not dominated by any other then there is some value $X/hour which you are implicitly placing on your time, so that you take all opportunities to spend an hour to make more than $X, and none which make less than $X, and you also take all opportunities to buy an hour for less than $X, and none for $X or more.[1]

2.2 Why we probably don't do terribly implicitly

Most people go about their day-to-day lives without ever considering trying to attach a numerical figure to the value of their time. We might therefore ask how far from optimally they tend to act. There are some reasons to expect this not to be terrible. For the larger the departure from rational behaviour, the easier it is to spot, and be confident in, a dominance argument. And if there were systematic large inefficiencies in the behaviours seen as societally normal, we could expect that people who did spot these and acted accordingly would tend to be more successful, which would then go some way towards normalising the deviant behaviours.

2.3 Why we probably don't do brilliantly implicitly

The argument above is in very broad terms, and doesn't bite sharply. It seems effective at saying that we shouldn't expect to be able to uncover efficiency gains of orders of magnitude. But for efficiency gains of the order of, say, 10%, it seems to have no weight. Other factors contribute to much larger differences in efficiency between people, so even if there is a correlation between people who are less suboptimal and those who do well, it will be hard to see the effect because it will be drowned out. Moreover at this level it could be much harder to make a compelling

[1] Strictly speaking there is another assumption going into this conclusion: that the opportunities are small, or divisible, so the effect of any opportunity on the marginal value of an hour or dollar is negligible. It could fail if they come in large blocky amounts. In fact employment opportunities often do come in such blocks, so the simple model at least needs refinement in some cases. Nonetheless it holds approximately for a range of activities.

argument out of dominance reasoning, because of uncertainty in the all-things-considered prices of different outcomes.

However a 10% increase in efficiency through your life could be a large deal to an individual, and well worth pursuing. So we're interested in understanding the shape of these tradeoffs, and how to weigh them appropriately.

2.4 Not all hours are equal: the role of motivation and energy

Of course considering opportunities as simple time/money tradeoffs is a simplification. To some extent this is a feature; we are seeking simpler models in order to give usable approximations that can aid decisions. But a poor model could easily be worse than no model at all, if it persuades us to follow reasoning to a counterintuitive conclusion. It can easily happen that the reason the conclusion seems counterintuitive is that it is in fact false, because of some factor not properly accounted for in the model. I think a simple time/money model is such an oversimplification.

The problem with the model is that it simply treats time as a resource to be spent. Any hour doing something uses up time from the same pool, and any hour gained adds to that pool. This ignores differences between different activities. This could be valid if pure time were the main restriction on what people could achieve. However for many people time is just one resource. Just as important are having the motivation to get on with the task at hand, and the mental energy to execute it well.

2.5 Examples: why it matters

Consider the case where you're deciding whether to take a taxi home, which will take 30 minutes (including waiting for it) and cost $20. The alternative is to walk, which would take 90 minutes. So you have an opportunity to spend $20 to save an hour. If you work flexible hours and are able to take on an extra hour's work for $30 after tax, it seems that there is a clear dominance argument that you would be better to take the taxi and work for the saved hour rather than walk home: you take the same amount of time and end up $10 richer. (It might be better still to take the taxi and not work for the saved hour.)

Perhaps, however, you enjoy walking. We can imagine that you would prefer to spend 90 minutes walking than that same 90 minutes split: 30 minutes in a taxi, 30 minutes working, and 30 minutes leisure. Now there's a new line of dominance reasoning available to us. From a baseline where you take the taxi rather than walking, assuming that you will spend some time working and some time in leisure, you would get an improvement by turning the 30 minutes of the taxi ride and any marginal half hour block of work and of leisure into the 90 minute walk: you'll end up $5 richer and also enjoy yourself more.

Clearly if we have two purported dominance arguments pointing in opposite directions, something has come off the rails. What? In this case it looks like the first dominance argument went wrong because it wasn't accounting for all of the relevant resources. It was indeed better in time and money terms to take the taxi (you could take the taxi and work for 50 minutes to end up better on both axes), but it was achieving this by spending happiness.

2.6 Why it's really more than one bucket

Our example makes the point that we need to consider more than just the tradeoffs between a time bucket and money bucket. We should at least account for some mental effects like the change to happiness, motivation, or mental energy levels. Now, could these just be different labels for the same thing? Some examples show that they cannot. For example if you play chess as a hobby it is quite possible that playing in a tournament would consume mental energy (perhaps more so than working), but increase your happiness and perhaps longer-term motivation.

2.7 Why to model it as one bucket for simplicity

Despite knowing that there are several different factors other than time and money, we may not want to model them all. The more variables there are in our model the harder it will be to use. We've made a case that just counting time and money is too few, but we're going to try to roll these other factors into one for consideration.

This is a point of our ensuing model which should be watched closely. It is possible that there are good reasons to count more of these separately, and you'd need to adjust the methods if this turns out to be the case.

2.8 Lack of natural units

One reason to be wary of trying to account for any of these things is that they are hard to measure, and there is no obvious natural unit to use. This probably explains why they're often ignored. As we'll see in the next section, our approach will be to bypass this to use a natural conversion rate to turn them into units we can more easily conceive.

2.9 Sleep

As well as the various different components of mental preparedness, physical tiredness can be an important factor in these decisions. This is of course complicated by the fact that there is an interaction between mental and physical tiredness. Physical tiredness itself has multiple components. One of these is a need to sleep. As this need is fairly uniform day-by-day, we're going to set it aside in the subsequent analysis. Every extra factor we model makes the model harder to use in everyday life, so we'd like to keep it simple where possible.

3. Neutral time costs

3.1 Factoring out the problem of units using intuited conversion rate

The approach we suggest for dealing with the problem of units and measurements is to convert by replacing the simple question:

> *How many hours would activity X take to complete?*

With the more nuanced question:

> *How many hours would you need to complete activity X and recuperate so that you feel as fresh at the end as at the start?*

The idea is that the relevant figure is usually not simply how long something takes to do, but should include the necessary recuperation. If calculated in these terms, we refer to the resulting answer as the *neutral time cost* (NTC). Its units are *neutral hours* (Nh).

Some activities are of course net refreshing, in which case no sensible answer to the above question exists. The idea for how to extend the definition is based on the following observation:

> *For any period of time T and activity X, doing X and recuperating in a period of length T + NTC(X) is as refreshing as spending T recuperating.*

This is since for activity X, doing X and recuperating in a period NTC(X) is by definition as refreshing as spending no time recuperating, so adding a period of T to each should preserve the equality. This suggests the definition:

> *For any time period T, activity X, let N(X,T) be the amount of time needed to complete activity X and recuperate so that at the end you are refreshed as you would have been spending period T recuperating. Then NTC(X) = N(X,T) – T.*

Although in theory this number is independent of choice of T, it's not clear that the estimates we'd provide would be independent of T. For example when T is very large we might think the activity X is almost irrelevant. To avoid ambiguity here, we suggest it is canonical to take the smallest usable value of T. If the activity is net tiring, this will usually mean T=0, recovering our first definition of NTC. If the activity is net refreshing, this generally means that T is the length of time equally refreshing as completing activity X (without any extra recuperation).

Also as remarked above we have many different mental resources, so "as refreshed as" isn't exactly defined. Nevertheless for practical purposes we are often able to give estimates for this, as we can easily see how to trade off between similar resources.

3.2 Consequences and subtleties of the definition

There are some cases where an activity draws in significantly different amounts on the different mental buckets we discussed earlier. In this case we might use a larger value of T, to allow the extra time to normalise out between the buckets. To return to our earlier example, it might be hard to assign a neutral time cost directly to participation in a chess tournament, because it is simultaneously draining and enjoyable, increasing longer-term motivation. So there is no length of recuperation that is equivalent. But perhaps playing in the tournament and recuperating for three hours is more balanced, and equivalently refreshing as eight hours of leisure. In that case the NTC of playing in the tournament would be the time taken up by the tournament, minus five hours.

Note that there is, essentially by definition, always an activity available of zero NTC. This is simply whatever you would be doing to recuperate. Could there be an activity of negative NTC? We might think not: if there were, *that* would become the best recuperation activity, and therefore would have zero NTC. However, we'd like to restrict the default recuperation to be a default, which will normally mean that it doesn't involve spending any other resources (such as money). By spending those resources you could get something that was negative NTC, and you could compare the neutral hours gained to the cost to decide if it were worth it.

An example of something with negative NTC might be buying (and drinking) a cold drink; this could easily be more refreshing than a rest of the same time period.

It looks like we may have to be careful about accounting for this "doesn't spend any other resources". For example if your current recuperation activity is reading, and you are reading the best book you have found all year, then your current recuperation probably has negative NTC. You are using up the resource of "book unread", as by reading now you mean you will have lower quality leisure time in the future.

3.3 Neutral Time Cost Ratios and Work Ratios

We have defined the neutral time cost of a given action. Sometimes that action might be "spend an hour doing Y". In this case we naturally have a ratio, the *neutral time cost ratio* (NTCR(Y)) = NTC(hour of Y)/hour. The units of NTCR are Nh/h, but if you identify neutral hours and hours you can regard this as a dimensionless ratio. It tells you how to convert, for example, from a $/h value for Y to a $/Nh value.

4. Worked examples

Here is a list of examples of how you might use a figure for $/Nh to make quick practical decisions. For all of the following we'll use a figure of 20$/Nh.

4.1 Taking work

Alice is offered a tutoring job at 50$ for an hour once a week. That works out as 40$ after tax. She'd need to do about 20 minutes prep work, which she estimates has a NTCR of 2.5, so a NTC of 50 minutes. She quite enjoys the tutoring, but it is somewhat tiring, and she estimates a NTCR of 1.5, and the hour has a NTC of 90 minutes. Finally it's a 15 minute walk each way, which is somewhat refreshing, so she says a NTC of 10 minutes each way. That's a total NTC of 50 + 90 + 20 = 160 minutes, so it works out as about 15$/Nh. That's below the threshold, so she decides against the job.

She's offered a second job at $60 for an hour, and only 6 minutes' walk away. This time the after tax pay is $48, and the total NTC is 50 + 90 + 8 = 148 minutes. That gives a price of $19.5/Nh. This is very close to her threshold, so she decides that it's worth a little more thought. On reflection she thinks she can make the preparation fun, and therefore a lower NTC, and she ends up taking the job.

4.2 Buying shirts

Bob is deciding between two shirts. He likes the first a bit more, and it is 20$ cheaper, but the second is non-crease.

He estimates that for the patterns he'd be willing to walk for about half an hour to change the less nice of the two patterns to the nicer one, so it's about 20 neutral minutes better (using the same estimate for walking as in the previous example). He expects to wear the shirt perhaps 50 times. Ironing a shirt takes him about 3 minutes, or 5 neutral minutes. So the non iron shirt is 5*50 = 250 neutral minutes better in expected ironing time. Since 250 neutral minutes is a lot better than 20 neutral minutes plus $20, he takes the non-crease shirt.

4.3 Choosing restaurant

A group of friends is deciding between restaurants A and B. Restaurant A is right next to them, and serves better food, but it's about $15/head pricier. Restaurant B is a 10 minute walk each way, but cheaper.

Walking with friends is much lower NTC than walking alone – it's nearly as good as general social time. They think the 20 minutes of total walking is a 5 minute NTC. In fact, though, they can shortcut this calculation: you just ask whether you'd rather have dinner at the nice and adjacent restaurant, but have to cut the evening 45

minutes shorter (as that's what $15 buys at $20/Nh). They decide the extra social time is more valuable than the extra food, so prefer B.

4.4 Travelling to a meeting

Claire and Daniel are colleagues. Claire has a hour meeting in another city, which is two hour's drive away. Daniel is wondering whether to travel with her. He thinks it is worth about 0.5Nh for him to be present at the meeting as well (not counting the time for the meeting itself) – which looks much less than the 4 hours required for the round trip.

However they realise that driving alone has a NTCR of about 1, driving in company has a NTCR of 0.4 as they can use the time productively in conversation. As a result, having Daniel accompany Claire costs Daniel 1.6Nh, but *gains* Claire 2.4Nh. They value their time at similar levels, so Daniel decides to go with Claire to the meeting.

5. Neutral time costs and energy levels

5.1 NTC varies with current energy level – appropriate to use correct costs

The definition of NTC is in terms of the time needed to return to the same state you began the work in. This is therefore really a function not just of activity, but of (activity, initial state) pairs. While it would be nice if the function was pretty independent of initial state, we need to be able to account for cases where they come apart. If measuring and comparing NTCs of different activities, the appropriate comparison to make is where each activity has its NTC calculated for the initial state at which the activity is being contemplated.

5.2 Extremes of energy give higher NTCs

Some thought experiments are useful here. Let's consider how NTC for a fixed activity (let's say writing reports) varies with initial level of mental energy. Firstly note that at very low energy levels, you'd typically find writing very hard and make very low progress per unit time. So the NTC must be high (indeed for low enough energy levels you would be unable to write at all without first resting).

What about very high levels of mental energy? A very high level of mental energy should correspond to a state where you feel motivated and clear-thinking. It is a state which would be efficient in the sense of the amount of writing you can do per unit time (also true for pretty much any other activity). However it is also a hard state to recapture, so the recuperation time needed to return to it would be high. Hence we also have a high NTC at very high energy levels.

That deals with the asymptotics. In the middle range I'd expect a fairly simple curve (though I don't have a strong justification for this, and could be wrong) and smooth behaviour. So the shape might be something like this:

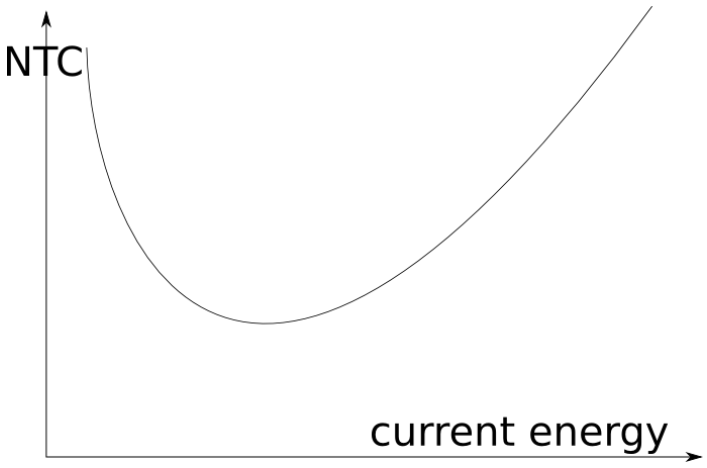

There will be a graph of this type for every activity.

5.3 Best to work at or near minimum NTC for that work

If we are looking to get the most done, we should be able to do that if we work at the energy level that minimises the NTC for the work. If the smoothness assumption holds the graph will be fairly flat for a region around this minimum, so there will be a range of energy levels which are all almost as good at the optimum (see figure below).

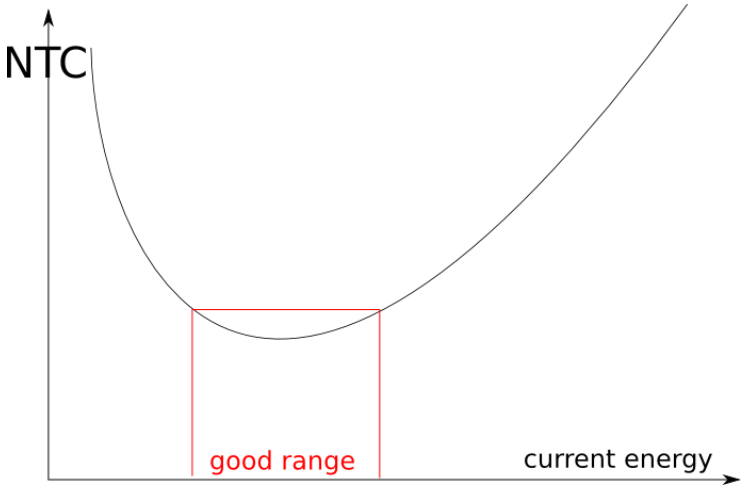

One type of behaviour we'd expect to see for some activities is that there is a region where they are net refreshing, and outside of that region they are net tiring. For example perhaps walking has this sort of behaviour: if you are very tired, walking will tire you; if you are very fresh, walking will take dull the edge, but there is a middle ground where it is recuperating.

5.4 Focus-demanding activities will have higher optima

Roughly speaking the shape of this graph is determined by two factors: how much recuperation time you need for a given period spent on the activity (NCTR), and how much you can achieve per hour. The first of these dominates at the right hand end of the graph creating the asymptotic there, and the second gives us the asymptotic on the left.

My guess is that there is rather more variation between activities in the second factor than the first, and that this is therefore the chief determinant in where the minimum lies. Some activities demand a lot of focus, and you may achieve almost nothing if you attempt them in a lower energy state. Considering important decisions falls into this category: the baseline for comparison is a quick decision, and unless you are sharp it's unclear extra time in consideration is valuable. For menial tasks, on the other hand, the output won't vary too much between initial energy levels.

5.5 Sometimes desirable to change energy level

In the simple model we've just presented you always want to perform activities at the optimum point on the curve. An assumption, not yet explicitly discussed, which we used to get to that conclusion, is that you want to recuperate so that you are as fresh at the end of a period as at the start.

If you want in the same day to spend some time on multiple activities, and these have different optimal energy levels, there's a question of how to proceed even in the ideal case. To move from a lower level activity to a higher level one is fairly simple: you just need to rest. To move from a higher level activity to a lower level one is trickier to understand. If the good regions in the energy scale overlap, you can just switch from one to the other at a point in the overlap. If not, you will spend some time running down between the regions on one of the two activities.

6. Applications to life

The idea of this section is to discuss how one might apply the theory to help decision-making.

6.1 Look for ways to decrease the NTC of frequent activities

A big win, if you can, is to decrease the NTC of things you have to do a lot. If something is more mental stress than it needs to be, then getting to a point where you're calmer about it could be very valuable. Similarly if you can make your recuperation more effective you could easily get substantial gains.

6.2 Try out a figure for $/Nh and see how it goes

On the question of how to value time, Tyler Cowen suggests picking a figure and living with it for a while to see how it goes. That is, you evaluate opportunities in terms of their implicit conversion between time and money, and take the ones which lie on the right side of your threshold. This may not be the best figure, but the idea is that in the first place it should allow you to eliminate dominated behaviour, and secondly that after experiencing it for a while you'll be better positioned to work out if it's too high or too low.

I think this is a good suggestion, but would change it to a suggestion that you use a figure for $/Nh instead of $/h. This has an added complication in that producing estimates of the neutral hours something will take is a bit harder than the hours it will take, but it has some major advantages. Firstly, since it's accounting for more resources, it's less likely to break and have you take a dominated course of action. Secondly, it's broader in scope, as it can more naturally compare changes to leisure activities to time-saving changes.

6.3 Evaluate purchases in terms of Nh

There are many things we buy which just improve our leisure time. For example books, games, going to concerts. How can we understand the gains here in terms of neutral hours? Asking questions like the following can be quite effective:

Would I be more refreshed by 30 minutes reading this book, or X minutes spent on other leisure?

If the value of X which makes you indifferent is, say, 40 minutes, and the book will take 6 hours to read, then the six hours reading the book will be as refreshing as 8 other leisure hours, so the book is worth 2 neutral hours to you.

You might prefer to ask in terms of what you would most enjoy (using this as a reasonable and easier-to-think-about proxy) rather than which would be most refreshing.

6.4 Caveat: don't trust calculations blindly

As always when using a theoretical model to inform practical decisions, don't be too constrained by the outputs of the model. If you get an answer that seems unintuitive, try to understand why it's unintuitive. If you can find a simple reason why you'd

expect your intuitions to be wrong here, the calculation may be correct. Otherwise it's quite likely that there's an important factor your model hasn't accounted for, but which your intuition has accounted for.

7. Questions

I've presented Neutral Hours as a model that we can use to reduce the highly multidimensional space of mental resources down to a single dimension to help practical decision-making. This model occupies a strange place between theory and practice. Of course this dimensionality reduction is crudely information-destroying, and wouldn't be endorsed by a complete theory of how to weigh mental resources. On the other hand there is too much detail and too much model to think about for many practical decisions. My hope is that it provides a theory which while imperfect is *applicable*: that we can use it not to make all of our decisions but to help us find better decision-making rules. Even for this purpose there may be many improvements that can be made.

In this section I'll present the questions I think it would be most useful to answer to help make this an even more usefully applicable theory.

7.1 What are the NTCs for different activities?

One obstacle to using these ideas is the difficulty in estimating NTCs for different activities. It could be useful both to come up with recommendations for how people can do this for themselves, and to have standard estimates for different activities that people can use (particularly where the variation between people is small), or that can be used in cost-benefit analysis.

7.2 What Nh/$ rates are offered by common time-saving opportunities?

How much do you need to value your time before you should get a dishwasher? When is it right to take a taxi rather than a bus? Doing some of the work for these common situations could help people to apply the takeaway lessons.

7.3 How should we count other resources?

Leisure may have other benefits than recuperation; it may also use other resources. How could we pull these into the framework while keeping estimation practical?