

The optimality of perception and cognition:

The perception-cognition gap explored

Andreas Jarvstad

Cardiff University

February 2012

Supervisors:

Prof. Ulrike Hahn

Dr. Simon K. Rushton

Dr. Paul A. Warren

Thesis submitted for the degree of Doctor of Philosophy

Summary

The ability to choose wisely is crucial for our survival. Yet, the received wisdom has been that humans choose irrationally and sub-optimally. This conclusion is largely based on studies in which participants are asked to make choices on the basis of explicit numerical information. Lately, our ability to make such high-level choices has been contrasted with our ability to make low-level (perceptual or perceptuo-motor) choices. Remarkably, we seem able to make near-optimal low-level choices. Taken at face value, the discrepancy gives rise to a *perception-cognition gap*. The gap implies, for example, that our ancestors were much better at choosing where to put their feet on a rocky ridge (a perceptuo-motor task), compared to choosing which prey to hunt (a cognitive task). The work reported herein probes this gap. There are many differences between literatures showing optimal and sub-optimal performance. The main approach taken here was to match low- and high-level tasks as closely as possible to eliminate such differences. When this is done one finds very little evidence for a perception-cognition gap. Moreover, once the standards of performance assessment of the respective literature are applied to data generated under such conditions it becomes apparent that the cause of the gap seems to lie in the standards themselves. When low-level standards are applied, human choice, whether low- or high-level, looks good. When high-level standards are applied, human choice, whether low- or high-level, looks rather poor. It is easy to see then, that applying high-level standards to high-level tasks, and low-level standards to low-level tasks, will give rise to a “gap”, with no or little actual difference in performance.

Acknowledgements

My heartfelt thanks

To Ulrike, Simon and Paul

for guidance, dedication and lively discussions.

Many a lesser intellect would have given up on this tenacious candidate many a time.

To Nadine (and our little baby “sessa” Freyja)

for patience, support and understanding.

Many a woman would have complained (more loudly).

To Ann and Christer

for being good parents (with all that this entails).

To those at Cardiff who in any way helped this project along

The Sensational Seminar Group, the Cog Group and (the now Dr) Jon Kennedy deserve special mention.

To Iain and David

for understanding that sometimes things don't quite work out as planned.

Thank you!

Table of contents

1.	General Introduction	7
1.1	The normative status of decision theory	8
1.2	The universal applicability of decision theory	11
1.3	Shortcuts to (good) decisions	12
1.4	Tasks underlying the perception-cognition gap	14
1.4.1	Cognitive (high-level) tasks	14
1.4.2	Perceptual and perceptuo-motor (low-level) tasks.....	16
1.5	Differences between low- and high-level paradigms	19
1.6	The perception-cognition gap explored	20
2.	Are Perceptuo-Motor Decisions Really More Optimal Than Cognitive Decisions?	23
2.1	Experimental investigation	24
2.2	Methods.....	27
2.2.1	Participants and Instructions	27
2.2.2	Apparatus	27
2.2.3	Stimuli, Experimental Design and Procedure	27
2.2.4	Data analysis	29
2.3	Results.....	31
2.3.1	Response time	31
2.3.2	Movement variability	32
2.3.3	Aim point error.....	34
2.3.4	Target choice behaviour.....	35
2.3.5	Task performance.....	37
2.4	Discussion – Experiment 1 and 2.....	39
2.5	The effect of task parameters on performance metrics	40
2.5.1	Task parameters and the optimality of aim point choices	40
2.5.2	Task parameters and the optimality of target choices	42
2.5.3	Task parameters and bounds of optimal performance for whole experiments....	44
2.6	General discussion	47
2.6.1	Summary of empirical results	47
2.6.2	Apparent relaxation of precision criteria when aiming for large targets.....	47
2.6.3	Apparent under-utilisation of response time	48
2.7	Theoretical implications.....	49
2.7.1	The Problem of Task Difficulty	49
2.7.2	On the use of optimal agents to infer unconditional optimality.	50
2.7.3	General implications for inferences about the optimality of behaviour	50

2.8	Summary	52
3.	Interlude 1	53
4.	Knowing When to Move On: Cognitive and Perceptual Decisions in Time	55
4.1	Experimental paradigm	56
4.2	Timing decisions based on perceptual information	57
4.2.1	Experiment 1: Motion discrimination & task difficulty	58
4.2.2	Experiment 2: Motion discrimination & changes in rewards and penalties.....	60
4.3	Timing decisions based on perceptual and cognitive information compared	62
4.3.1	Experiment 3 – Motion discrimination, mental arithmetic & mental rotation....	62
4.3.2	Experiment 4 – Dynamic changes in task difficulty	64
4.4	General discussion	65
5.	Interlude 2	67
6.	The Gaps That Weren't: The Perception-Cognition & The Description-Experience Gap .	70
6.1	Methods.....	73
6.1.1	Participants and apparatus.....	73
6.1.2	Stimuli, design & procedure	73
6.2	Results & discussion	76
6.2.1	How good were people's choices?.....	76
6.2.2	Do people deviate systematically from optimality?.....	79
6.2.3	Are peoples subjective beliefs calibrated?	81
6.2.4	Do people deviate from optimality when biases are taken into account?	83
6.3	Summary	85
7.	General Discussion	88
7.1	Summary of empirical work	88
8.	Supplemental materials.....	92
8.1	Chapter 2 – Supplementary materials	92
8.1.1	Movement & reaction time analyses.....	92
8.1.2	Control experiment	94
8.2	Chapter 4 – Supplementary materials	96
8.2.1	SI Methods: Fitting the Weibull function to accuracy data	96
8.2.2	SI Methods: Assessing goodness of fit	97
8.2.3	SI Methods: Additional details	98
8.2.4	SI Methods: Mathematical formulation of the decision problem	100
8.2.5	SI Methods: Extrapolation from Weibull fits to choice data	101
8.2.6	SI Methods: Was the hard-easy manipulation successful?	102
8.2.7	Fig. S4.1.....	104
8.2.8	Fig. S4.2.....	105

8.2.9	Fig. S4.3	106
8.2.10	Fig. S4.4	107
8.2.11	Fig. S4.5	108
8.3	Chapter 6 – Supplementary materials	109
8.3.1	SI Additional methods: Model fitting	109
8.3.2	Fig. S6.1	114
8.3.3	Fig. S6.2	115
8.3.4	Fig. S6.3	116
8.3.5	Fig. S6.4	117
8.3.6	Fig. S6.5	118
8.3.7	Fig. S6.6	119
9.	References	120

1. General Introduction

Many, if not all, of our actions cause outcomes only probabilistically. This implies that we cannot decide between possible actions solely on the basis of their values. Instead, we have to take both the likelihood and the values of outcomes into account. But how should we trade-off values and likelihoods when choosing amongst actions? The optimal strategy is to choose the action with the highest expected (subjective) value (Bernoulli, 1738/1954, Von Neumann & Morgenstern, 1944/1955, Savage, 1954/1972). The expected value of an outcome is the product of its likelihood and its value. Choosing optimally, therefore, entails choosing the action associated with the largest value-probability product.

Decades of research studying human high-level cognitive decisions suggest that humans deviate from this optimal choice strategy (Allais, 1952/1979, Ellsberg, 1961, Kahneman & Tversky, 1979; Tversky & Kahneman, 1981; 1992). Such deviations have been found mainly (but not only, see e.g., Fox & Tversky, 1998) using decision tasks in which participants are asked to choose between options, for which probabilities and values are given in numerical format. As a participant you might, for example, be asked to indicate whether you prefer option A: “£4000 with a probability of .8”, or option B: “£3000 with certainty” (Kahneman & Tversky, 1979). Asking such questions researchers have been able to show that people choose in a manner incompatible with the optimal strategy (see “Cognitive [high level] tasks” below).

When the same normative standards are applied to low-level (perceptuo-motor and perceptual) decisions, however, they appear to describe the observed choices very well. So well, in fact, that peoples choices can be described as optimal, or near-optimal (e.g., Trommershäuser, Maloney & Landy, 2003a, 2003b; Whitely & Sahani, 2008; Navalpakkam, Koch & Perona, 2009, Navalpakkam, Koch, Rangel & Perona, 2010). Thus, there exists a *perception-cognition gap*: low-level decisions appear optimal and high-level decisions appear sub-optimal (Trommershäuser, Landy & Maloney, 2006; Maloney, Trommershäuser & Landy, 2007; Trommershäuser, Maloney & Landy, 2008).

The main difference between low- and high-level studies is that in the former participants are not asked to choose between options with probability information in a numerical format. Instead, they are asked to choose between actions, for which the probability of success and failure is derived from low-level systems (see “Perceptual, perceptuo-motor (low-level) tasks” below).

The perception-cognition gap implies, for example, that our ancestors were much better at choosing where to put their feet on a rocky ridge (a perceptuo-motor task), compared to choosing which prey to hunt (a cognitive task). Why should this be? If one is puzzled by this gap, a natural starting point is to try to equate as far as possible low- and high-level decisions in order to compare them fairly. This was the main approach taken in the work presented here.

We begin by discussing briefly general questions that a reader may have at this point: Why is expected (subjective) value maximization optimal? How can one think of the perceptual system as making decisions? What exactly are cognitive and perceptual decisions? How have researchers determined whether their participants adhere to the optimal strategy? And so on...

1.1 The normative status of decision theory

Why is choosing the option with the highest expected (subjective) value considered to be the optimal strategy? There are two main arguments for its normative status: axiomatic and long-run performance arguments. The long run argument is perhaps the most intuitive. An agent who chooses the option that has the highest average value will, by definition, do at least as well *on average* as other agents employing *any* other strategy. So, if you want to do as well as conceivably possible across your lifetime, you should make choices that maximize expected value. One problem with this formulation is that it is not clear that it generalises to the case when you choose only once (or sufficiently infrequently for the law of large numbers to apply) – specifically, the fact that the expected value of a particular action is positive will not console the person who has just lost everything (Jensen, 1967).

Axiomatic developments do not rely on the law of large numbers, but start from a set of rules. From such rules, or axioms, a decision theoretic framework prescribing how a person who wishes to abide by these axioms should act can be developed. There exist many axiomatic developments (e.g., Von Neumann & Morgenstern, 1955; Savage, 1972). Introductions to decision theory tend to describe 3-4 axioms (e.g., Jensen, 1967; Berger, 1985, Parmigiani & Inoue, 2009).

Berger's (1985) description includes four axioms. Axiom 1 requires that one must have a full set of preferences: either you prefer A to B, B to A or you are indifferent between A and B. In other words, for any two options you either prefer one or the other, or you are indifferent between the two. Axiom 2 requires that your preferences are transitive: if you prefer A to B, and B to C, you will also prefer A to C. Put differently;

if you prefer a banana to an apple and an apple to a pear you will prefer a banana to a pear. Axiom 3 states that if you prefer A to B then you prefer a probabilistic mixture of A to the same probabilistic mixture of B. To be precise, for any probability p you prefer $pA+(1-p)C$, where C is a new option, to $pB+(1-p)C$. That is, adding C to B and adding C to A will not change your preferences. Axiom 4 states that there are no infinitely bad or good outcomes: if you prefer A to B and B to C, there will be some probabilities (α , β) such that B is preferred to $\alpha A + (1-\alpha)C$, and such that $\beta A + (1-\beta)C$ is preferred to B. If you behave in accord with these axioms you will be maximizing expected utility, but why should you behave according to the axioms? That is, why is it rational to maximize expected utility in the manner prescribed by the axioms?

Typically, justifications for behaving as the axioms prescribe can be given by appeal to desirability. For example, if your preferences are not transitive (Axiom 2 above) you can be turned into a money pump. You may prefer a banana to an apple, an apple to a pear, but a pear to a banana. With those preferences, we can sell you an apple for a pear and a little money (e.g., £.10). We can then sell you a banana for the apple we have just sold you and a little money. Because you prefer a pear to a banana we can now give you back your pear and receive a little money plus the banana. If we repeat this procedure you would soon be without money.

Nevertheless, the normative status of decision theory has been challenged. Some critiques are aimed at the normative status of decision theory itself. Searle (2001), for example, argues that he would never wager his life against twenty-five cents, no matter how small the likelihood of death (i.e., he does not approve of Axiom 4 above). Searle's argument may seem compelling at first. However it is worth putting the example into the perspective of everyday life. Just by getting out of bed, driving a car, exercising, cooking food, in short by living ones ordinary life one constantly exposes oneself to substantial risks (that have low potential payoffs).

Do you really need to go on holiday for example? No? Are there additional risks associated with going on holiday that you would otherwise not be exposed to? Remember it does not matter how small those risks are – as long as your life is at risk. Searle's argument implies that you would never go on holiday (as e.g., air travel is associated with an additional small risk of death) – yet many people do go on holiday, suggesting that Searle's argument cannot be quite right.

A few challenges pertain to whether decision theory is a *reasonable* standard for human behaviour. These are typically based around the idea that being “optimal” requires full knowledge and that the theory quickly becomes intractable under such

conditions (e.g., Gigerenzer, 2008). We do not consider such arguments particularly damaging to the theory's use here for three reasons. Firstly, one does not need to assume that someone who acts according to decision theory is an all-seeing all-powerful agent. One simply needs to assume that the agent chooses according to decision theory on the basis of the information that is available to it in its environment (i.e., is statistically optimal). Secondly, issues of computational tractability are often bound to specific algorithmic implementations. Even if decision theory were computationally intractable, given a specific task with a specific knowledge set, it might still be approximated by heuristics. More importantly, in the current context, we do not consider issues of tractability relevant to decision theory's status as a normative framework, and therefore whether it can, or should, be applied to human behaviour. In other words, whether or not humans can be expected to be able to choose the best action is orthogonal to the issue of whether they *do* choose the best action (or not).

We note that decision theory is a normative theory that might not apply directly in some specific contexts. It is, for example, only normative for so called games against nature. Games against nature are choice situations in which one does not face an adversary – hence against nature. When an adversary is present, and there is competition for resources, the normative theory needs to be extended to take into account the fact that actions of one agent might affect other agents (von Neumann & Morgenstern, 1953).

Similarly, standard decision theory typically assumes that utilities are state-independent. This means that your preferences for certain outcomes are independent from the state that you are in. For example, if one prefers cappuccino to tea then strictly speaking one always prefers cappuccino to tea. However, one might, for example, imagine a scenario in which one's preference for caffeine-containing drinks was related to caffeine blood-levels, such that for higher blood-levels, drinks lower in caffeine are preferred. In the same way that decision theory has been extended to account for competitive choice, it has also been extended to allow for state-dependent utilities (Arrow, 1973; Karni, Schmeidler & Vind, 1983).

Although both the mentioned extensions seem entirely reasonable they make the theory more complex. Because the perception-cognition gap arises in situations which do not require state-dependent utilities or competition amongst decision makers (compare e.g., Trommershäuser et al., 2003a and Kahneman & Tversky, 1979), we do not deal with either here. Moreover, up until this point, we have used “subjective” in brackets, when discussing value maximization, to denote the fact that the standard

framework allows subjective values (utilities) to differ from objective values. This aspect of decision theory was introduced to account for diminishing marginal utility of money and has been given normative grounding (e.g., Bernoulli, 1738/1954).

Diminishing marginal utility is intuitive; £100 when one has £100 000 in the bank will be perceived as less valuable compared to £100 when one has £10 in the bank. Here, however, we will restrict ourselves to expected value maximization as outlined next.

Recently, Rabin (Rabin, 2000; Rabin & Thaler, 2001) has argued convincingly that value preferences should be linear over quite a large range of values in order for the theory plausibly to claim normative status. Here we will be concerned with values considerably lower than those considered by Rabin, and will be concerned with quantities that have monotonic relationships with real money (points in experimental tasks).

Moreover, the low-level decision literature has typically used expected value maximization as a normative standard. As it is a stricter criterion of optimality than expected utility maximization, and it is generally used in low-level studies (e.g., Trommershäuser et al, 2003a, Whitley & Sahani, 2008), it seems an appropriate benchmark reference when comparing across low- and high-level decision tasks. For these reasons, we will from this point onwards use expected *objective* value maximization as the normative standard. This also simplifies estimation of optimal strategies considerably as one does not have to estimate peoples' value weighting function.

1.2 The universal applicability of decision theory

When people think about decision-making, they are likely to think about high-level, effortful and conscious choices. They might think, for example, about deciding which of several houses to purchase or about deciding how to invest savings. However, decision theory can be applied much more broadly. In principle, it can be applied to any system or organism that can “select” amongst two or more actions.¹

Anderson (1990), for example, shows that human categorization, human memory, and human problem solving can be thought of as solving the problem of minimizing a

¹ We eschew here the tricky philosophical issues of free will and determinism. We assume merely that for most actions organisms in general, and humans especially, have two or more actions from which they can “select”. Whether this selection process is in some sense “free” or fully determined is for the purposes of determining whether or not a system behaves in accord with decision theory irrelevant. For example, an animal may have in its behavioural repertoire fight and flight responses. The animal's ability to engage fight and flight responses in exactly the right situations is orthogonal to whether these behavioural responses are in some sense “free” or pre-determined (since the beginning of time).

cost function that takes into account both uncertainty and values. Furthermore, the ability of organisms to make choices has been studied for a wide range of behaviours and organisms, from the foraging behaviour of birds (Pompilio & Kacelnik, 2010) to the paths amoeba “choose” to take (Nakagaki, Yamada & Toth, 2000). In short, anything that offers an organism or a system more than one “option” can be analysed from a decision theoretic perspective.

In fact, it might be argued that decision theory *should* be applied this broadly. An example in point is the traditional distinction in psychology between judgment and decision-making (e.g., Feldman, 2006). The reason for this distinction is presumably that judgments are viewed as passive estimates about some property of the world, whereas decisions are assumed to involve choosing among possible future states of the world (perhaps on the basis of underlying judgments). However, the idea that judgments are passive can be criticised. It can be argued that judgments generally are *for* something. That wider purpose will determine the cost function (see e.g., Berger, 1985, and see Harris, Corner & Hahn, 2009 for empirical evidence that people are sensitive to this). For example, if you need to judge the width of a stream in order to jump over it, underestimating the width is more costly than overestimating it. On the other hand, if you and a friend compete about who can provide the most accurate width estimate, underestimation and overestimation are equally costly.

In other words, systems or organisms generally perform actions for reasons. Reasons determine, jointly with the environment, the cost function. Neglecting this seems to at best be harmless (e.g., when the implicitly assumed cost function is the correct one), and at worst may produce misleading results (e.g., when the implicitly assumed cost function does not match the one participants have).

1.3 Shortcuts to (good) decisions

It was noted that decisions involve taking into account both probabilities and values. Above, we suggested that any system or organism could, and perhaps should, be analysed from a decision theoretic perspective. This may seem odd. It might seem preposterous to suppose that, for example, amoebae have access to separate estimates of values and probabilities and can combine them as decision theory suggest that they should. The application of decision theory to these simple organisms might therefore be viewed as misguided.

An alternative view is that the mere application of decision theory to study a system does not also imply a particular algorithmic implementation of decision theory

(in that system). Instead, decision theory can be used as a normative standard. As such, it can be used to chart the efficiency of systems and organism regardless of the algorithms and mechanisms that underlie their behaviour.

In fact, in many situations there is no need for an organism to represent values and probabilities separately. Or, for that matter, to weight explicitly values and probabilities - regardless of whether the organism is a human or a rat. If an animal is given the opportunity to experience outcomes, it can instead learn which of several options it prefers.

There are many such reinforcement learning algorithms (see e.g., Sutton & Barto, 1998), which although not proven to converge, tend to converge on the optimal solution (in the sense that they maximize expected reward) given sufficient experience. That is, an organism can be well described by decision theory despite not explicitly operating according to its principles (i.e., combining separate estimates of value and probability multiplicatively).

It seems important therefore to distinguish decisions that are made in environments where such learning is possible from decisions that de facto have to be based on estimates of probabilities and values - because no such learning could reasonably have taken place. The ruling out of learning strategies is particularly important in the current context where we wish to compare high-level decisions to low-level decisions. Classical high-level cognitive studies of decision-making do not typically allow participants the possibility to take such short-cuts.

In classical cognitive studies, participants instead receive descriptions of probabilities and values and have to make hypothetical decisions on the basis of these descriptions. They receive no feedback and cannot therefore learn which option is the better one. To pick the optimal option, they *have* to combine the values and probabilities in an optimal manner. Studies of perceptual and perceptuo-motor decisions, however, typically do provide feedback. A trivial explanation for the perception-cognition gap, therefore, is that performance in low-level studies is good because people are given feedback and poor in high-level studies because they are not.

Can optimal perceptuo-motor, or perceptual, decisions be explained away as a result of a simple learning process? Some studies explicitly model learning (e.g., Navalpakkam et al., 2009). The authors of these studies are presumably not worried about ruling out learning strategies, but apply decision theory in the same way it can be applied to rats or amoebae. However, many recognize the challenges posed by alternative explanations based on learning algorithms. Whitley & Sahani (2008), for

example, offered their participants only intermittent feedback to minimize the possibility that participants might gradually home in on an optimum. Similarly, Trommershäuser et al. (2003a) argued that learning cannot explain their results as A) the individual data shows no gradual improvement across time and B) when the task changed, participants seemed to immediately find the new optima and did not have to go through an exploratory phase of gradual improvement (as expected from a system that learns). The latter method was later developed as a more general test of optimal performance (Bayesian transfer, Maloney & Mamassian, 2009). Nevertheless, learning may be rapid and learnt solutions may generalise across conditions. Thus, if one wants to compare performance for low-level decisions to performance in paradigms where learning is not possible (classical decision tasks) it seems prudent to eliminate feedback entirely.

In the first study presented here, we explore perceptuo-motor performance and therefore provide feedback as is typical in these paradigms (see e.g., Trommershäuser et al., 2003a,b). In later studies, however, we avoid feedback to ensure that we are studying decision-making in the classical sense (and not the learning of the setting of decision criteria through feedback).

1.4 Tasks underlying the perception-cognition gap

1.4.1 Cognitive (high-level) tasks

Since its inception, expected value/utility theory has been tested by asking people to consider hypothetical choices between options. A common format is to offer the choice between two options, A1 and B1 (from Kahneman & Tversky, 1979):

Table 1.1. Example of a standard choice problem in the classical literature.

A1				B1			
<i>£</i>	<i>p</i>	EV	EU	<i>£</i>	<i>p</i>	EV	EU
4000	0.80	3200	51	3000	1	3000	55

As a participant, you receive only the information in the first two columns under each option: the monetary value of each option (£) and the probability of the realisation of the monetary value (*p*).

If you are like most people, you would choose B1. As the EV column shows this choice violates expected value theory ($EV(A1)=3200 > EV(B1)=3000$). It does not,

however, violate expected utility theory. If you are allowed to weight your values, by say using a logarithmic value function (e.g., $v(x) = x^a$, where $a = .5$), you are considered optimal when choosing this option ($EU(A1)=51 < EU(B1)=55$).

However, demonstrations of irrational choice in the classical literature do not typically rely on such single choices, but on juxtaposing options, such as the one just presented, with another set of options (for a more recent example of this approach see Birnbaum, 2008).

Table 1.2. Example of a standard pair of choice problems in the classical literature.

A1				B1			
<i>£</i>	<i>p</i>	EV	EU	<i>£</i>	<i>p</i>	EV	EU
4000	0.80	3200	51	3000	1	3000	55
A2				B2			
<i>£</i>	<i>p</i>	EV	EU	<i>£</i>	<i>p</i>	EV	EU
4000	0.20	800	13	3000	0.25	750	14

The second option-pair – A2 and B2 – is illustrated in Table 1.2 above. If you are like most people, you would choose A2 here. On its own, this choice is consistent with expected value maximization. However, if we apply the same weighting we applied previously in order to capture people’s choices for the first pair (A1 & B1), we see that you should have chosen B2. Thus, your choice pattern is inconsistent – both when evaluated by the expected value and by the expected utility maximization norms. If you behaved as either an expected value or an expected utility maximizer you could not have these preferences. In terms of axiomatic expected utility theory your choice pattern violates the independence axiom (Axiom 3 above).

Using this method of juxtaposing choices across pairs of options, researchers have developed many combinations of options for which people’s preference patterns violate axioms of expected utility theory. The implied logic is as follows: if it can be shown that people have preferences that systematically violate expected utility theory they cannot be said to make choices in accord with the same theory - and are therefore sub-optimal.

Note that classical studies employing the above paradigm have mainly used between-subject methodology, have not provided feedback, and have not involved real payoffs (e.g., Kahneman & Tversky, 1979). We view the absence of feedback as an

important strength of this paradigm. Nevertheless, when feedback is provided, performance often improves (e.g., Chu & Chu, 1990; Shanks, Tunney & McCarthy 2002; Jessup, Bishara & Busemeyer, 2008). This suggests that people can learn to make less sub-optimal choices if given the opportunity to do so. In other words, it suggests that really irrational behaviour might be less likely to persist in the real world where feedback is often available (but for an example of apparently irrational decision-making in the presence of feedback see the decision from experience literature, e.g., Hertwig & Erev, 2009).

Classical studies have been criticized for the lack of real payoffs. As a participant, you are asked to choose the option you *would* prefer. However, your choices are of no consequence to you. As a relatively disinterested participant, you may not be willing to invest the mental effort required to make wise choices (e.g., Smith, 1976) on the basis of inconsequential word problems. Recently, it has even been suggested that effort should be explicitly incorporated in theories of choice (Dickhaut, Rustichini & Smith, 2009). This critique presupposes that if sufficient cognitive effort were induced, for example by making choices consequential, decisions would no longer show deviations from optimality. The evidence for better performance with increased payoffs however is mixed (see e.g., Camerer & Hogarth, 1999; Hertwig & Ortmann, 2001 for reviews). In brief, although some deviations from optimality may be avoided when actions are consequential, it is far from clear that all or even most deviations cease to exist for (large) real potential payoffs.

1.4.2 Perceptual and perceptuo-motor (low-level) tasks

As noted initially, the main difference between low- and high-level decision tasks is that participants have to use their lower-level systems (perceptual and perceptuo-motor) in the former to inform their decision making. However, this is not the only way in which the decision making paradigms differ. The low-level studies that have been used to contrast people's ability to make low-level decisions with their ability to make high-level decisions (e.g., Trommershäuser et al., 2006) also differ from high-level studies in how adherence to optimal standards is assessed and in the method by which decision-making is studied.

Fig. 1.1 illustrates a typical low-level paradigm. Because the perceptuo-motor system is noisy, speeded pointing towards a target will result in responses dispersed around the chosen aim point (cross, Panel A, Fig. 1.1). In Trommershäuser et al.'s (2003a, 2003b) paradigm, participants point under time pressure towards stimulus

configurations (Panel B) with the goal of earnings as many points as possible. Participants accrue points if they hit a reward region (full line, Panel B), lose points if they hit a penalty region (dashed line, Panel B), and receive both if they hit the intersection of both regions. Different aim points (different symbols, Panel B) will result in different probabilities of hitting each region (*hit probabilities*, Panel C). Different hit probabilities, in turn, will result in different number of points earned.

Given that there are many aim points, participants are in effect choosing between many different options of the form: reward with $p = X$, penalty with $p = Y$, both reward and penalty with $p = Z$ – which is easily recognized as a classical decision-making problem (see e.g., Tversky & Kahneman, 1979). As is evident however, the choice in Trommershäuser et al.'s paradigm (2003a) is not a binary one – between two choice options – but a choice along a continuous scale: the x and y position of the aim point. To fulfil the task goal of earning as many points as possible, participants have to choose the aim point with the highest expected value.

Note that participants have to use knowledge derived through the perceptuo-motor system in order to make their decisions. That is, the probability information is not provided as numbers on a piece of paper as in the classical tasks, but must be derived from low-level systems. In this example, participants have to assess how likely they are to hit each region given specific aim points. Note also that the value information, however, is given in abstract form just like in the classical paradigm.

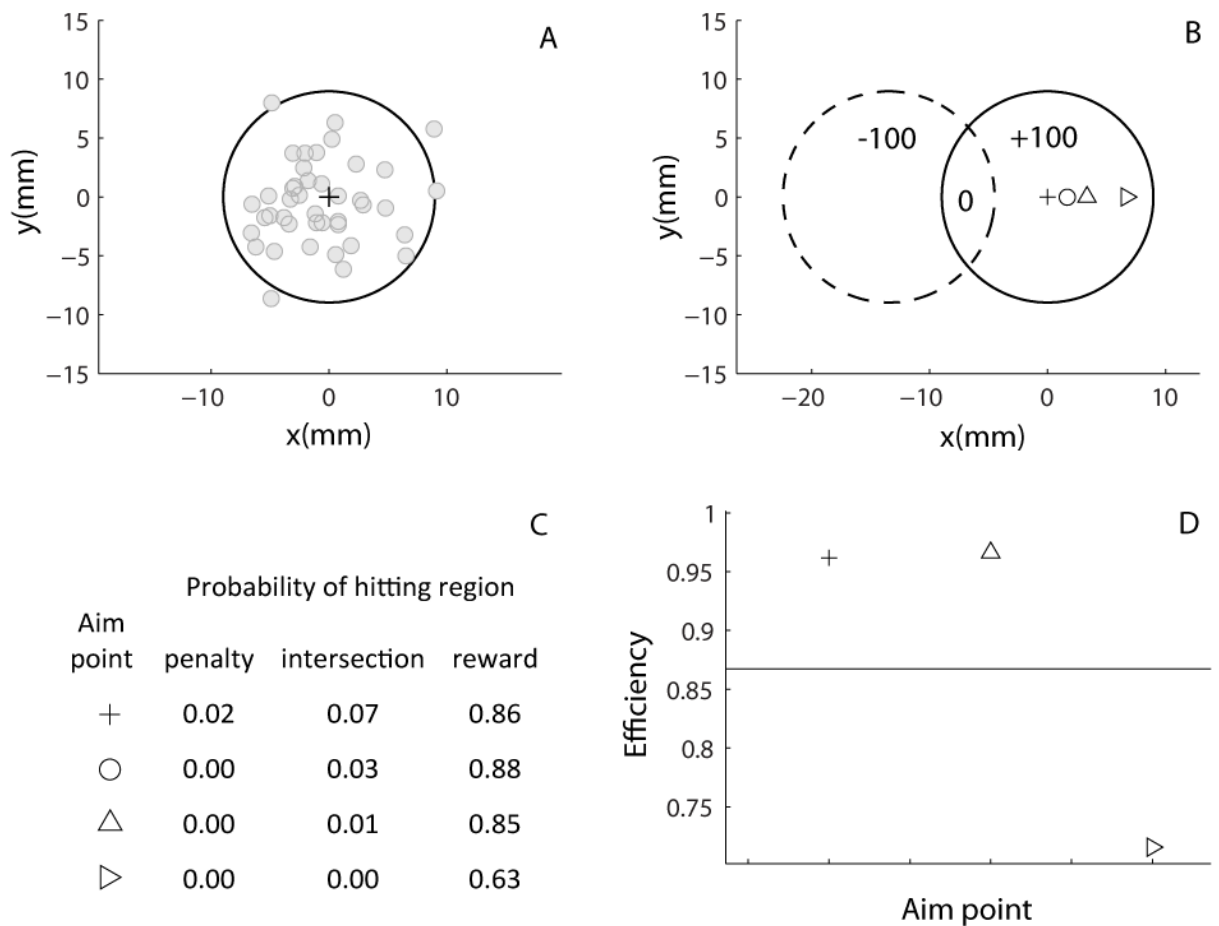


Fig. 1.1. Perceptuo-motor gambles and performance assessment. Panel A: A simulated response distribution (grey discs) from one participant ($\sigma^2 = 14.78$, Participant 2, Exp. 2, in Trommershäuser et al., 2003a) aiming at the centre of a target (cross, Panel A). Panel B: Example of one stimulus configuration and reward structure employed by Trommershäuser et al., with example aim points (symbols) and region-specific rewards and penalties (numbers). Panel C: Hit probabilities for the aim points in panel B. Panel D: Efficiencies (expected gain normalized by optimal expected gain) for the aim points in panel B. The optimal aim point (circle, Panel B), has an efficiency of 1. The horizontal line represents the lower 95 percentile of optimal performance. Efficiencies below this line are lower than expected by chance and hence sub-optimal.

To assess participants' performance, in paradigms such as the one just outlined, researchers typically use ideal observer methods. The idea is to compare participants' performance to the performance of a hypothetical participant who performs the task optimally. Performing the task optimally means to, given the available knowledge, make the best possible choices. The best aim-point choice - the choice that returns the maximum amount of points (and therefore money) - is the disc in Panel B.

Participants' choices can be compared to the optimal ones. Consider, for example, the upwards-facing triangle in Panel B, which represents a participant who is a little too careful and aims further away from the penalty region than optimal. This aim point will on average return a sizable fraction of the optimal agent's earnings. Panel D shows the different efficiencies associated with the various aim points in Panel B. Efficiency is a given participants' earnings proportional to the earnings of the corresponding optimal agent; an efficiency of 1 means that the participant is precisely optimal and efficiencies below 1 means that they perform worse than optimal.

Of course, given a limited sample size, and given the noise in the perceptuo-motor system, even an optimal agent is unlikely to achieve an efficiency of 1 for a particular experiment. To take this noise into account and to evaluate statistically whether particular choices can be said to deviate from the optimal ones bootstrap/Monte Carlo methods are used (e.g., Trommershäuser et al., 2003a).

Because a model of how the task should be performed is available (the ideal agent), optimal agents performing a particular experiment many times over can be simulated. Such simulations lead to distributions of optimal earnings. These distributions can then be used to infer whether a particular participant's earnings are significantly different from that of a hypothetical participant performing the experiment optimally. Typically, the lower 2.5 percentile of the optimal earnings is taken as a cut-off for optimal performance (see e.g., Trommershäuser, Gepshtein, Maloney, Landy & Banks, 2005). This lower threshold for characterizing participants as optimal is illustrated in Fig 1.1 D by the horizontal line. In our example two of the three aim points (cross & upwards-facing triangle, Panel D) result in efficiencies above the lower 2.5 percentile of the optimal efficiency and are thus classed as optimal.

1.5 Differences between low- and high-level paradigms

The just outlined paradigm is different from the classical cognitive paradigm in that it offers the possibility of learning the optimal strategy (through feedback), offers many repeated choices between the same options (one target-reward configuration is considered as one option), typically offers more than two choice options and offers consequential choice (participant payment is contingent upon performance). Any of these differences, individually or jointly, could potentially explain the perception-cognition gap. In the following, we will seek to minimize these differences by studying perceptuo-motor, perceptual and cognitive choices under precisely matched conditions.

The low level approach also differs in important ways from the classical paradigm with regard to how performance is assessed. As discussed above, the classical paradigm uses carefully tailored sets of choices, across which particular choice patterns indicate violations of fundamental axioms of decision theory. By contrast, perceptuo-motor studies compare actual earnings to the earnings achieved by someone who chooses in accord with decision theory. Consequently the perception-cognition gap could also be due to differences in how performance is assessed. For example, people's good performance in low-level tasks may be a result of them using choice strategies that only approximate the optimal ones. If the approximation is sufficiently good, however, then performance will not be classed as statistically different from optimal. This would be an interesting result in itself, and would suggest that the view of human cognition as severely flawed (e.g., Sutherland, 2007) is an exaggeration. Indeed, such a result might be taken to suggest that violations shown in cognitive studies are relatively harmless.

In the first two chapters we use the low-level approach of evaluating actual performance exclusively. However, in the third chapter, we use both types of performance evaluation across precisely matched low- and high-level tasks. This allowed us to empirically evaluate the idea that the gap might be due to the application of different performance standards.

1.6 The perception-cognition gap explored

We began our exploration of the perception-cognition gap with a study of perceptuo-motor decision-making (Chapter 2). Our task was based on a standard perceptuo-motor decision-making paradigm (Trommershäuser et al., 2003a, 2003b). Under time pressure, participants tried to hit targets by pointing at them. They received feedback (points like in computer games) and were paid as a function of how well they did. We manipulated target distance and target size in a first exploration of these factors in the perceptuo-motor decision literature.

As typical in perceptuo-motor decision studies, participants had to choose an aim point for each stimulus configuration – an “implicit” choice. We added a more cognitive “explicit” choice to the task. On each trial participants were presented with two targets: a small and a large one. Thus, in addition to deciding where to aim on a particular target, they had to decide which of the two targets to aim for. To decide well, participants had to take into account their own perceptuo-motor uncertainty, the distance to each target and the size of each target (all three factors determining the likelihood

with which targets can be hit) together with the rewards of hitting the targets (the small target was always worth more than the large).

We initially thought that adding a more “cognitive” component to a perceptuo-motor task might make perceptuo-motor decisions more like those in the cognitive domain – sub-optimal. We did not find quite what we expected. The results *did* indicate that participants did not optimize two performance-related metrics (precision and time usage). More importantly from the perspective of the overall goal here: simulations and comparisons across our studies demonstrate that optimality depends on task difficulty. Thus, the standard analysis employed in perceptuo-motor decision-making experiments seemingly fails to provide an absolute standard of performance. It is therefore unclear how different domains can be compared. This, in conjunction with non-trivial evaluative and methodological differences, was a first indication that comparative claims favouring perceptuo-motor, or perceptual, systems over higher-level cognitive systems might be premature.

In the absence of an absolute standard of performance, we sought to find way in which we could compare performance across domains without confounding performance with task difficulty. That is, we wanted to rule out, as far as possible, the possibility that differences in performance were due to, for example, cognitive tasks being more “difficult” than perceptual tasks. The idea behind our next set of experiments (Chapter 4) was to use a decision task that might be viewed as modality independent, thus minimizing the potential for the chosen decision problem to be more or less suited to a particular system. One such task is making decisions about how much time to spend.

To decide how much time to spend on a given task wisely, you need to know how costly it is to get the task wrong, how rewarding it is to get it right, and how your task performance changes as a function of how much time you spend on the task. Importantly, it is a type of decision that we make for low- as well as high-level tasks.

We first investigated timing decisions when the underlying task was perceptual. Decisions were highly efficient and suggested that people can make good use of perceptual knowledge and abstract reward information. We then compared timing decisions for the perceptual task to timing decisions for more cognitive tasks. Performance was highly similar, suggesting that knowledge can be acquired, and used to make timing decisions, in an equally efficient way regardless of whether the knowledge is derived through perceptual or cognitive experience.

Although the results of Chapter 4 showed that cognition can be as good as perception they left open the possibility that the equal performance shown was due to the domain in which participants made decisions. Decision about time might be special. Moreover, although we used two cognitive tasks neither included numerical probability information as is typical in classical cognitive studies. The next line of work we present (Chapter 6) was designed to address these two potential issues.

The study reported on in Chapter 6 also allowed us to address another decision making gap: the *description-experience gap* (Hertwig & Erev, 2009; Rakow & Newell, 2010). In classical tasks with numerical probabilities and values participants *overweight* low probabilities. In contrast, when participants can learn values and probabilities low probabilities are *underweighted*: the *description-experience gap*. However, as for the perception-cognition gap, confounds between tasks makes comparisons difficult.

Avoiding the typical confounds, we compared choices across three precisely matched tasks: a classical decision task with numerical information and two tasks for which the numerical probability information was replaced with equivalent low-level (perceptuo-motor) and high-level (mental arithmetic) information. Comparisons across the three tasks suggests A) that the perception-cognition gap is illusory and due to differences in how performance is assessed, B) that the description-experience gap is due to the assumption that objective probabilities match subjective ones (and/or due to learning in decision from experience studies), C) that deviations from optimality observed in classical decision-making studies might not be particularly costly and finally D) that individual differences are more important for predicting peoples' choices than the type of decision people face.

2. Are Perceptuo-Motor Decisions Really More Optimal Than Cognitive Decisions?²

As noted in *General Introduction*, there appears to be a striking dissociation between human perceptuo-motor- and cognitive decision-making performance. Normative decision theory poorly describes cognitive decision-making (Birnbbaum, 2008; Kahneman, Slovic, Tversky, 1982; Kahneman & Tversky, 1979). Perceptuo-motor decision-making, on the other hand, appears well described by the same theory (for a review see Trommershäuser, Maloney & Landy, 2008; see Whiteley & Sahani, 2008 for a similar conclusion in a perceptual domain). This apparent dissociation has been highlighted repeatedly. Trommershäuser, Landy and Maloney, for example, note that "... in marked contrast to the grossly sub-optimal performance of human subjects in traditional economic decision-making experiments, our subjects' performance was often indistinguishable from optimal." (2006, p. 987; see also e.g., Maloney, Trommershäuser & Landy, 2007; Trommershäuser et al., 2008).

This performance dissociation is puzzling. Few reasons are evident for why perceptuo-motor decision-making should be optimal, while cognitive decision-making is sub-optimal (but see e.g., Chater & Oaksford, 2008; Evans & Over, 1996). Furthermore, little progress appears to have been made in explaining the difference.

There are at least three possible sources for the apparent dissociation: 1) competence may be modality dependent 2) performance may be task dependent and 3) differences may result from the way performance is evaluated. If competence were indeed modality dependent this would be a striking finding. However, as pointed out by Trommershäuser and colleagues (e.g., Maloney et al., 2007), the employed experimental paradigms differ along a number of methodological dimensions. Perceptuo-motor studies generally involve repeated decisions with outcome feedback and internalized probabilities. Cognitive decision tasks, on the other hand, generally involve one-shot decisions without feedback and exact probabilities stated on paper (see e.g., Birnbbaum, 2008; Kahneman & Tversky, 1979, but see e.g., Hertwig, Barron, Weber & Erev, 2004; Thaler, Tversky, Kahneman & Schwartz, 1997). Thus, a less interesting explanation is that one, or many, of these methodological differences give rise to the apparent dissociation.

Not only are there methodological differences, performance is also evaluated differently in the two fields. Although both perceptuo-motor and cognitive studies draw

² A version of this chapter is under review in *Cognition*. A pilot study with a similar experimental design was submitted for partial fulfilment of a Master's degree in Research Methods.

on normative theories to provide performance standards, adherence to these norms is assessed in different ways. Generally, the perceptual, and perceptuo-motor, literature asks how closely human performance matches that of an ideal agent (see e.g., Barlow, 1962; Geisler, 2003; Trommershäuser et al., 2003a, 2003b). Broadly, an ideal agent is a model that performs a given task maximally well. Constraints under which the system is assumed to operate are typically built into the model. The cognitive literature, on the other hand, typically asks if a system violates some (or many) of the principles of normative theories (e.g., Birnbaum, 2008; Hertwig, et al., 2004; Kahneman & Tversky, 1979). Experiments are designed so that certain response patterns will violate fundamental axioms of decision theory. Thus, assessment of performance differs in two ways: absence³ versus presence of system constraints and qualitative versus quantitative violations of normative theories (but see Wu, Delgado & Maloney, 2009 for an attempt at equating tasks across domains).

Given the just outlined non-trivial differences between cognitive and perceptuo-motor studies, comparisons of human performance across the two domains need to be made with care. Here we highlight difficulties associated with such comparisons using two perceptuo-motor decision-making experiments. The experiments demonstrate that minor changes in task parameters, changes which do not impact on an optimal participant's performance, influence whether participants are viewed as optimal or sub-optimal. We follow up these empirical results by illustrating, through simulations, how specific changes in task parameters can cause participants hitherto classified as optimal to be classed as sub-optimal. The experiments also suggest that people's perceptuo-motor decisions are sub-optimal in ways not captured by Trommershäuser et al.'s (2003a, 2003b) model. Together these results, we think, suggest that claims of greater optimality for perceptual systems over higher-level cognitive systems may be premature.

2.1 Experimental investigation

Using the perceptuo-motor paradigm outlined in *Chapter 1: Perceptual, perceptuo-motor (low level) tasks*, or variants thereof, Trommershäuser, Maloney and

³ Studies of higher-level decision-making and judgment typically are not concerned with constraints when evaluating participant performance. Instead it is assumed that the paradigm employed is sufficiently easy, so that any system that adheres to the studied axioms is able to perform the necessary computations (Evans, 1993). This is not to say that constraints have gone unstudied. Kahneman and Tversky (1996), for example, have argued that when extensional cues are given to participants, performance improves. This effect is presumed due to extensional cues triggering a slow and effortful processing system that would otherwise not have been used (Kahneman & Frederick, 2002).

Landy have explored perceptuo-motor decision-making extensively (see Trommershäuser et al., 2008). We were initially interested in one of the distinctions they make: that of implicit and explicit decisions. Seydell, McCann, Trommershäuser and Knill (2008) note that cognitive paradigms generally involve explicit choices (introspectively one is aware of choosing), whilst perceptuo-motor paradigms generally involve implicit decisions (introspectively one is unaware of choosing). Trommershäuser et al. have previously explored the explicit/implicit choice dimension in two studies (Trommershäuser et al., 2006; Seydell et al., 2008) – and concluded that explicit as well as implicit motor choice is optimal, or near-optimal.

In our experiments (illustrated in Fig. 2.1, see Methods below for details) designed to explore this distinction further, participants made two choices per trial: an aim point choice (“implicit”) and a target choice (“explicit”). All pointing movements originated from a dock (white disc) and targets were displayed at different distances. On each trial, participants had to choose whether to attempt to hit a small or a large target. Hitting a target incurred a reward (the small target was always worth more than the large target) and missing a target incurred a penalty. The task goal was to earn as many points as possible. To earn as many points as possible, participants had to trade off the probability of hitting each target with its associated values. Target hit probabilities depended on participants aim point choices, their motor variability, the size of the target, and the distance to the target.




Fig 2.1. Design of Experiment 1 & 2. Crosses indicate potential target locations. Grey discs represent one possible target configuration. The white disc represents the dock from which all movements originated. The reward for hitting the large target was 50 in Experiment 1 and 75 in Experiment 2. In both experiments, the reward for the small target was 100 points, and the penalty for missing either target was -25. Note: targets are not drawn to scale.

A novel aspect of our study was that the expected gain (the number of points you would expect to receive on average when trying to hit a target) of each target depended on the size of the target as well as its distance to the dock (Fitts, 1954; Schmidt, Zelaznik, Hawkins, Frank & Quinn, 1979). Thus, a basic question was whether humans are able to trade off these quantities in an optimal manner when making perceptuo-motor choices.

The use of two target sizes also enabled an indirect assessment of one the assumptions built into Trommershäuser et al.'s model (2003a, 2003b), namely the assumption that motor error is unconditionally minimized. This assumption is critical for previous studies for two reasons. Firstly, it is a basic building block of Trommershäuser et al.'s (2003a) model and other models of motor planning (e.g., Harris & Wolpert, 1998). Secondly, it is important as the assumption that motor error is minimized is carried through to the modelling of optimal choice and hence the evaluation of participant performance.

Previous studies have also probed the question of human time allocation in perceptuo-motor tasks. The general conclusion has, again, been that time allocation is optimal or near-optimal (e.g., Battaglia & Schrater, 2007; Dean, Wu & Maloney, 2007; Hudson, Maloney & Landy, 2008). However, in these past studies participants were explicitly instructed to optimize time usage. Consequently, this does not answer the question of whether the perceptuo-motor system optimizes time in general.

Relevant to the latter issue is the study of Gepshtein et al. (2007). This study employed near and far targets and a fixed response time and it found that participants reached faster to near targets than to far targets – even when the same amount of response time was available for both distances. In other words, participants did not maximize time use for near targets. The speed-accuracy trade-off (Fitts, 1954; Schmidt et al., 1979) describes the inverse relationship between pointing precision and movement speed: the faster the movement the lower the precision. Given this trade-off, it appears that participants' failure to maximize time usage for near targets resulted in decrease in precision and potentially a decrease in rewards obtained. As Gepshtein et al.'s results suggest that time allocation in motor responding may not be optimal without specific, explicit instruction, further examination seems important.

We conducted two experiments with the task just outlined. The task parameters differed across Experiment 1 and 2. Specifically, target size, target distance, number of possible target locations and the reward for the large target differed across the

experiments (see Methods for details and see Fig. 2.1 for an illustration of some of the differences). To state that the perceptuo-motor system is optimal (or nearly so), presumably implies that it can deal with a variety of situations that might occur – not that it is optimal for one particular target size or one reward structure only. That is, if the perceptuo-motor system is optimal, one would expect it to be able to cope with the changing conditions across Experiment 1 and 2. In the following, we report on both Experiment 1 and Experiment 2 simultaneously. This facilitates comparisons between the two experiments, which should produce very similar results. As it turns out, seemingly innocuous changes in task parameters can have dramatic effects on whether participants are classed as optimal or sub-optimal.

2.2 Methods

2.2.1 Participants and Instructions

Sixteen (8 in each experiment) members of the School's participant panel were paid an hourly rate of £6 to participate and received a performance related bonus based on their efficiency (efficiency * £6).

Participants were informed of the reward structure in each experiment and were told to maximize their total score (“earn as many points as possible”). Participants were told that they could receive an additional bonus of min £0 and max £6, the amount to be determined by their performance (“the better you do the more money you will receive”). All participants were naive as to the purpose of the study. All had normal, or corrected to normal, vision and were fully mobile. Participants were fully informed about the experimental protocol.

2.2.2 Apparatus

The experiments were written in Matlab (Mathworks, Inc.) and run with the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) on a Mac Mini (Apple, Inc.). Participants were seated in front of a pen display (Wacom DTZ-2100, Wacom Co. Ltd.) slanted at 65°. The pen display was used to display stimuli and record responses. Responses were recorded with the spring loaded eraser end of a standard Wacom stylus pen. Participants chose their distance and height relative to the display so as to enable natural pointing movements.

2.2.3 Stimuli, Experimental Design and Procedure

Fig. 2.2 (Panel A, see also Fig. 2.1) illustrates the possible stimulus configurations in Experiment 2 (however, the actual background used was black and not white). In both experiments, each stimulus configuration contained a dock (radius 16

pixels/~4.3mm) identifying the starting position. Two discs (potential targets), one large (Experiment 1: radius 16 pixels/~4.3 mm; Experiment 2: radius 22 pixels/~5.9 mm) and one small (Experiment 1: radius 8 pixels/~2.16 mm; Experiment 2: radius 11 pixels/~2.9 mm), were displayed to the left of the dock (except for one left handed participant, for whom dock/targets were mirrored).

In each trial, one disc was displayed on the 'up' axis and one was displayed on the 'down' axis. In Experiment 1, discs were displayed at one of two distances relative to the dock: near (200 pixels/~5.4cm) and far (900 pixels/~24.3 cm). In Experiment 2 discs were displayed at one of three distances: near (170 pixels/~4.6 cm), medium (340 pixels/~9.2 cm), or far (510 pixels/~13.8 cm). A full factorial combination of elevation, target location and non-target location resulted in eight unique perceptuo-motor stimulus configurations in Experiment 1 and 18 configurations in Experiment 2.

Each experiment consisted of one learning session (22 trials per target size and location combination) and one experimental session (44 trials per unique stimulus configuration). In the learning session, no explicit (target) choice was made. Instead, a disc was designated as the target by the colour green (the non-target was red), and participants simply had to hit the target disc. In the experimental session both discs were yellow and participants chose which of the two discs they wanted to aim for.

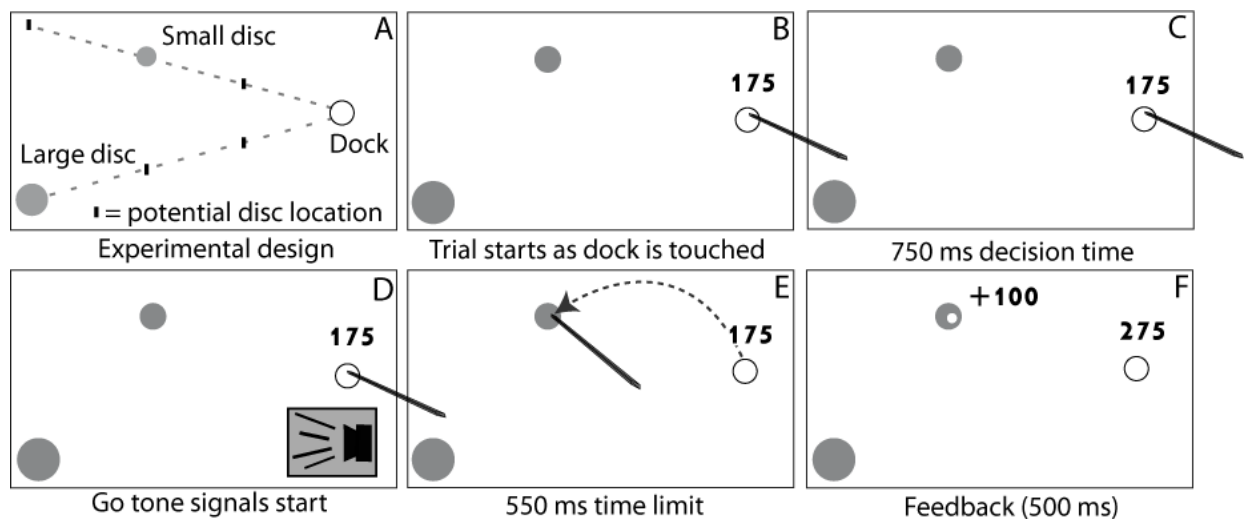


Fig. 2.2. Experimental procedure. Panel A: Possible stimulus configuration in Experiment 2. Panel B-F: Sequence of events as a trial unfolds. The number above the dock (white disc) represents participants' cumulative score. Note, stimuli are not drawn to scale.

In Experiment 1, the small target was worth 100 points, the large target was worth 50 points, and the background was worth -25 points. In Experiment 2, the reward associated with the large target was raised to 75 points, a manipulation that an optimal participant should be unaffected by.

Throughout the experiment, participants' cumulative score was displayed above the dock in blue numerals (Panel B-F, Fig. 2.2, – exemplified here by “175” and “275”). Participants initiated each trial by touching the dock with the stylus (Panel B), whereupon one of the unique stimulus configurations was displayed. Participants were required to maintain contact with the dock for 750 ms ('decision time', Panel C). A 550 Hz tone signalled that movement should begin (Panel D). After the tone, participants had 550 ms to attempt to hit their chosen target (Panel E). Participants received feedback both on where they hit the screen and on the amount of points earned on each trial (Panel F). They could rest at any time during the experiment simply by not initiating a new trial.

On a given trial, participants needed to respond within the 550 ms interval, but they were free to move as quickly as they wished within that upper bound. Responses that exceeded 550 ms were recorded as 'late'. Trials in which the stylus was lifted off the dock before 100 ms had passed since the 'go' signal were recorded as 'anticipatory'. Late and anticipatory responses resulted in feedback to speed up and slow down respectively and were rerun. The decision time and response time limits used match those of a previous study (Seydell et al., 2008).

For each trial, reaction time, movement time, response coordinates and points were recorded. Reaction time was defined as the time from the go-signal to the lifting of the stylus pen off the dock area. Movement time was defined as the time from lifting the stylus off the dock area to contact with the tablet surface. Total response time was the sum of reaction time and movement time. Response coordinates were defined as the x and y position of the stylus upon first contact with the screen after the stylus had been lifted off the dock.

2.2.4 Data analysis

The first block in the experimental session was treated as a warm up block and was deleted prior to any analyses. Late and anticipatory responses were discounted (see e.g., Seydell et al., 2008). For the decision session, the mean proportion of late responses was .07 (SD=.06). The mean proportion of anticipatory responses was .07 (SD=.05).

To assess participants' overall performance, a reliable estimate of movement variability is needed. The free choice component of the decision session meant that some targets (e.g., large near targets) had few or no data points. In order to guarantee a minimum of 20 data points for each estimate of movement variability the last 20 trials (for each target size and location combination) of the learning session were combined with the trials from the decision phase as the basis for estimates of participant's precision.

In deriving these estimates, outliers (defined as data points further than 2.5 times the large target radius from the target centre following Gepshtein et al., 2007) were excluded. The mean proportion of trials excluded as outliers in the merged data sets was .01 (SD=.016).

Responses were analysed separately for each participant and each factor (target size and target location). As in previous studies (e.g., Trommershäuser et al., 2003a, 2003b; Gepshtein et al., 2007) three assumptions were made. Firstly, it was assumed that the response distributions were bivariate normal, an assumption that was verified by inspecting chi square plots (Johnson & Wichern, 1998). Secondly, it was assumed that participants select a single aim point per target. In other words, it was assumed that the centroid of each response distribution describes the aim point for that distribution. Any deviation from this aim point was assumed to be due to unexplained variability influencing planning (Churchland, Afshar, & Shenoy, 2006) and execution (van Beers, Haggard & Wolpert, 2004) of movements. Finally, it was assumed that differences in biomechanical cost between targets were negligible (see e.g., Trommershäuser et al. 2003a, 2003b; Gepshtein et al. 2007).

To describe participants' pointing behaviour we use two metrics - aim point error and movement variability - which we computed separately for each participant's target size and target location combination. Given a normal response distribution, circular targets, and symmetric penalty regions (as employed here) the optimal aim point is the target centre. Aim point error describes the distance between participants' aim points (the centroid of each response distribution) and the target centre⁴. The lower than aim point error – the closer to optimal the aim point. Movement variability was defined as the mean distance of the movement end points from the centroid of the response distribution (see e.g., Gordon, Ghilardi & Ghez, 1994).⁵ Movement variability describes

⁴ Defining aim point as the [x, y] coordinate of a maximum likelihood fitted bivariate Gaussian (cf., Gepshtein et al., 2007) produced equivalent results.

⁵ Because movement data was anisotropic, defining movement variability as the standard deviation of the response distribution necessitates two dependent variables. Following Gordon et al. (1994) results in a

how variable participants' pointing movements were (their perceptuo-motor variability). For clarity of presentation target elevation was collapsed across when computing these two metrics and when describing participants' use of response time.

We present both individual plots as well as group averages for each analysis. Repeated measures ANOVA's were used to test for group-level effects. When sphericity assumptions were violated Greenhouse-Geisser corrections were used. Next, we report on how participants used the available response time. Thereafter we describe how movement variability and aim point choice relates to target distance and size. Following this, data describing participants' choices between the two targets (target choice) is presented. Finally, participants' overall task performance is compared to that of an optimal agent.

2.3 Results

2.3.1 Response time

Did participants use all of the available response time as in studies with only one effective reach distance (Trommershäuser et al., 2003), or did they fail to maximize time usage as in a previous study utilizing different reach distances (Gepshtein et al., 2007). As can be seen in Fig. 2.3, when targets were far away, participants used nearly all the available response time (550 ms).⁶ However, for near and medium distance targets participants used less than the available time (effect of target distance, Experiment 1: $F(1,7) = 85.14$, $p < .001$, $\eta_p^2 = .92$, Experiment 2: $F(2,8) = 247.15$, $p < .001$, $\eta_p^2 = .97$). This suggests that participants may be satisficing rather than maximizing time use. If they had used the maximal amount of time available there would be little difference between near and far targets and the plots in Fig. 2.3 would look like horizontal lines.

Another trend worth noting is that participants appear to use more of the available time when they reach towards small targets (dashed lines, Fig. 2.3) compared to when reaching to larger targets (full lines, Fig. 2.3). The difference between movement times for small and large targets was marginal in Experiment 1 ($F(1,7) = 4.87$, $p = .063$, $\eta_p^2 = .41$) and significant in Experiment 2 ($F(1,7) = 20.87$, $p = .003$, $\eta_p^2 = .75$). We did not detect an interaction between target size and target distance in Experiment 1 ($F(2,8) =$

univariate dependent measure, making analyses easier and the exposition clearer. Seydell et al. (2008) likewise adopted a univariate measure (the square root of the determinant of the covariance matrix) to describe the variability of anisotropic data.

⁶ Note that unless participants want to time out ~ 50% of the time, the mean response time has to be lower than the maximum response time.

2.17, $p = .184$, $\eta_p^2 = .24$), but did so in Experiment 2 ($F(2,8) = 5.92$, $p = .014$, $\eta_p^2 = .46$). For a detailed analysis breaking down the effects of response times into its separate components (reaction time and movement time) see Chapter 2 – Supplementary Materials.

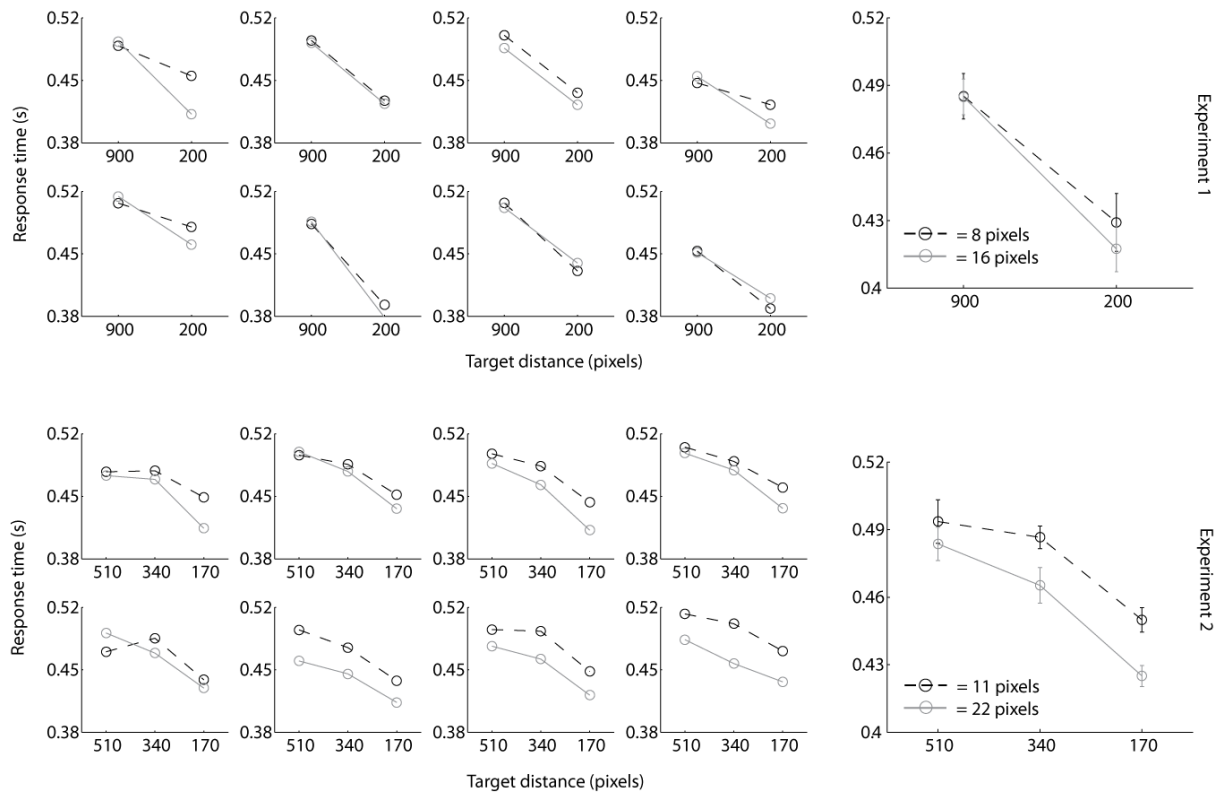


Fig 2.3. Response times: group averages and individual response time as a function of target distance, target size and experiment. The dashed line represents small targets and the full line represents large targets. The legend shows the radius of each target in pixels (1 pixel = .27 mm). Error bars are 95% confidence intervals useful for within-subject comparisons.

2.3.2 Movement variability




Fig. 2.4 Movement variability: group averages and individual movement variability as a function of target distance, target size and experiment. The dashed line represents small targets and the full line represents large targets. The legend shows the radius of each target in pixels (1 pixel = .27 mm). Error bars are 95% confidence intervals and facilitate within-subject comparisons.

Movement variability appears related both to target distance and size (Fig. 2.4). As expected, movements to far targets were more variable than movements to near targets (Experiment 1: $F(1,7) = 63.21$, $p < .001$, $\eta_p^2 = .9$, Experiment 2: $F(1.2, 8.4) = 40.9$, $p < .001$, $\eta_p^2 = .85$). Interestingly, movements were generally more variable for large targets (grey lines) than for small targets (dashed lines) in Experiment 2 ($F(1,7) = 15.47$, $p = .006$, $\eta_p^2 = .69$, size-distance interaction: $F(1.1, 7.7) = .23$, $p = .668$, $\eta_p^2 = .03$). In Experiment 1, this contrast was not significant ($F(1,7) = .78$, $p = .41$, $\eta_p^2 = .1$), but there was a marginal interaction between target size and distance ($F(1,7) = 5.23$, $p = .056$, $\eta_p^2 = .43$). If a direct statistical comparison between near small targets and near large targets is made (the likely origin of the marginal interaction), it reveals that movements to large near targets were more variable than those to near small targets ($t(7) = -4.14$, $p = .004$). Thus, in Experiment 1, participants aimed with greater precision to small near targets than they did to large near targets.⁷

⁷ There are trends in the data that suggest that for the furthest distance tested (Experiment 1, 900 pixels distance), the difference may disappear or even reverse (a trend that is also visible in the movement time plots, see Fig. S2.1). A possible explanation is that at very high difficulties participants relax their

A number of movement planning theories propose that movements are planned so as to minimize end-point variance (Harris & Wolpert, 1998; Trommershäuser et al., 2003a). That movements to larger targets are noisier than those to smaller targets suggests that the perceptuo-motor system does not always minimize end-point error but may instead adopt a satisficing approach (Simon, 1959). We will return to this issue and its wider implications below.

2.3.3 Aim point error

Aim point error is an indication of how well participants chose aim points (the implicit component). It describes how far participants aim points were from the target centre (the optimal aim point, see Methods-Data analysis). Compared to the highly consistent patterns for movement variability (Fig. 2.4), there appears to be little evidence for consistent between-subject patterns (Fig. 2.5). In other words, aim point choices do not seem strongly influenced by either target distance or size.




Fig. 2.5. Aim point error: group averages and individual aim point error as a function of target distance, target size and experiment. The dashed line represents small targets and the full line represents large targets. The legend shows the radius of each target in pixels (1pixel = .27 mm). Error bars are 95% confidence intervals.

precision criteria even further (e.g., “there is no point in trying hard – it’s too difficult”). An alternative explanation is that the far distance employed in Experiment 1 was sufficiently far, given the time deadline, as to constrain the possible pointing strategies that could be employed (i.e., it was not possible for subjects to choose different movement times for these targets).

In Experiment 1, there were no significant main effects of either size ($F(1, 7) = .02, p = .9, \eta_p^2 < .01$) or distance ($F(1,7) = 4.42, p = .074, \eta_p^2 = .39$), but there was a significant interaction between the two ($F(1, 7) = 12.18, p = .01, \eta_p^2 = .64$). In Experiment 2, there was a significant effect of size ($F(1,7) = 10.46, p = .014, \eta_p^2 = .60$), with aiming towards larger targets worse than aiming towards smaller targets. Inspecting individual data, this effect appears driven by some participants aiming more poorly towards nearer targets, and others aiming more poorly towards far targets, creating an overall effect of target size. There was no effect of distance: $F(2,14) = 2.45, p = .12, \eta_p^2 = .25$) nor was there a significant interaction ($F(2,14) = .86, p = .45, \eta_p^2 = .11$). Note, however, that aim points rarely deviated from the target centre by more than 5 pixels (1.35 mm), suggesting that participants' aiming performance was good.

2.3.4 Target choice behaviour

To describe participants' target choices, we compared the proportion of times the small target was chosen to the number of times it should have been chosen had participants been optimal. In Fig. 2.6, the proportion of small target choices is plotted as a function of the difference between the expected gain for the small and large target (ΔEV). If participants' choices were optimal, participants would always choose the small target for positive ΔEV (a small choice proportion of 1), and always choose the large target for negative ΔEV (a small choice proportion of 0). Cumulative Gaussians have been fit to the individual data to assist the eye. If participants were optimal, these functions would approximate step-functions centred on the dashed line at 0 ΔEV .




Fig. 2.6. Target choices. *Experiment 1 and 2*: plots show the proportion of times the small target was chosen (y-axis) as a function of the difference in expected value between small and large targets (ΔEV , x-axis), for each participant (each panel shows data from one participant). Each symbol (triangles in Exp. 1, discs in Exp. 2) represents one unique choice situation. For positive ΔEV the small target should be chosen (proportion small target choices should be 1) and for negative ΔEV the large target should be chosen (proportion small target choices should be 0). The lines are cumulative Gaussian density functions to facilitate comparisons across participants and experiments. The slopes of the functions are a measure of participant sensitivity to ΔEV . The intercept, or where the function intersects a small choice proportion of .5, is an indication of bias. Were participants' choices unbiased the intercept would be at or near $\Delta EV=0$. *Group level fits*: cumulative Gaussian density functions fit to data pooled across participants for Experiment 1 (full grey line) and Experiment 2 (dashed black

line) respectively. Error bars are bootstrapped 95 percentile intervals on the intercept estimate.

The individual data (Experiment 1 & 2, Fig. 2.6), suggest that participants are sensitive, but not perfectly so, to expected gain differences. Participants generally picked the higher valued target. However, differences between the experiments are apparent. In Experiment 1, many observers appear nearly un-biased. They choose small targets when these have higher EV's and large targets when these have higher EV's. In Experiment 2, on the other hand, most participants appear biased towards the small target – choosing it even if doing so results in a loss relative to choosing the larger target (the functions appear shifted to the left relative to 0).

To characterise this apparent bias on a group level, we pooled the data by experiment and fit cumulative Gaussian density functions. As can be seen (Fig. 2.6), group level fits confirm the apparent trend and show that the small target bias is stronger in Experiment 2 than in Experiment 1 (as judged by non-overlapping 95 percentile intervals). The relatively strong small target bias in Experiment 2 is noteworthy as participants appear to have aimed for the harder-to-hit target even though aiming for the easier-to-hit larger target would have resulted in a higher return.

2.3.5 Task performance

Task performance depended on two choices – choice of aim point and choice of target. An optimal agent always picks the best target and aim point.⁸ As the response distributions were Gaussian and the penalty region symmetric (i.e., missing either target incurred a penalty), the optimal aim point was always the centre of each target. For each participant, we simulated an optimal agent performing the experiment 100 000 times. The resulting distribution of average gains allowed us to estimate the expected gain of the optimal agent and the confidence in this estimate. If a participant's performance lay outside the lower 95% confidence bound they were classed as sub-optimal. If

⁸ It has been suggested that as a result of optimizing biomechanical cost the perceptuo-motor system is biased towards undershooting targets (Elliott, Helsen & Chua, 2001; Elliott, Hansen, Mendoza & Tremblay, 2004; Lyons, Hansen, Hurding & Elliott, 2006). Sometimes, undershoot refers to the spatial location (primary movement end point) of the initial (more or less) ballistic phase of movements (primary sub-movements, e.g., Lyons et al., 2001 p. 97). Sometimes, this type of undershoot refers only to movements that hit the target (% of undershoot/overshoot, Elliott et al., 2004, pp., 346-347). Since our apparatus did not allow for reliable trajectory measurements it is impossible to say whether primary end points undershot targets. However, we found some evidence of end point undershooting, but participants did not consistently undershoot the targets consistent with previous findings (Fitts & Petersen, 1964).

participants performed better than this lower bound, they were classed as statistically indistinguishable from optimal. In other words, we used standard methods to assess whether participants were optimal or not (see Trommershäuser et al., 2003a, 2003b for mathematical details).




Fig. 2.7. Overall task efficiency (white discs) and the lower bound of optimal efficiency (full line) for each participant in Experiment 1 (top panel) and 2 (bottom panel).

Fig. 2.7 shows participants' (x-axis) efficiencies for Experiment 1 (top panel) and Experiment 2 (bottom panel) respectively. The first thing to note is that participants' efficiencies are not distributed around an efficiency of 1 – as expected if participants were optimal. Nevertheless, six of eight participants in Experiment 1 were within the bounds of optimal performance. In Experiment 2, on the other hand, only one of eight participants' efficiencies was within the 95th percentile. A Fisher's exact test comparing the rates of optimal performance in Experiment 1 to Experiment 2 is significant ($p = .04$) and a Bayesian comparison of rates (Kass & Raftery, 1996; Lee & Wagenmakers,

2005) of optimal performance shows that the hypothesis that the rates of optimal performance (Exp 1, 6/8; Exp 2, 1/8) differ across the two experiments is 10.7 times more likely than the hypothesis that the rates are the same.

Nevertheless, the absolute efficiencies across the two experiments are fairly similar. That is, relative to the optimal agents, participants earned similar amounts in both experiments. The lower bound on optimal performance, however, appears to be substantially lower in Experiment 1 than in Experiment 2. Thus, the reason participants are classed as optimal in Experiment 1, and not in Experiment 2, appears to be due to differences in the confidence intervals not due to differences in absolute efficiency levels.

A Bayesian t-test comparing absolute efficiency levels across the two experiments shows that there is insufficient evidence to conclusively favour either the null or the alternative hypothesis (JZS Bayes Factor in favour of alternative hypothesis = .55, $t(14) = -1.13$, $p = .28$). However, the same test performed on the lower 95% confidence interval of optimal performance shows overwhelming support for the alternative hypothesis of a difference in confidence bounds (JZS Bayes Factor = 79438, $t(14) = -9.68$, $p < 1e-6$). We return to this issue below (“The effect of task parameters on performance metrics”).

2.4 Discussion – Experiment 1 and 2

Across Experiment 1 and Experiment 2, the experimental set-up was identical, and both experiments required two kinds of choices (aim point and target choices). However, the precise stimulus configurations and the reward structure differed across experiments. Compared to Experiment 2, Experiment 1 had smaller targets, fewer target locations, greater target-distance differences and the difference between the rewards for the small and the large target was greater.

It turns out that the differences in task parameters were highly consequential. Experiment 1 resulted in optimal participants, whereas Experiment 2 resulted in sub-optimal participants. This result implies that optimality standards as commonly employed are not absolute but relative. Relative standards imply that classifying systems as optimal, or sub-optimal, without further clarification is problematic. For which experiment should we use if we wanted to evaluate the optimality of the perceptuo-motor system: Experiment 1 or Experiment 2? We explore the effects task parameters have on the two sub-components of our task in greater detail below and return to this point in the General Discussion.

Regardless of whether participants were classed as sub-optimal or optimal, they were generally sensitive to the difference in expected gain between small and large targets. They generally choose the higher valued target more often than the lower valued target. On the other hand, participants in Experiment 2 were biased towards choosing the small target, representing the higher but more uncertain gain, *even* when this choice on average produced lower gains than choosing the lower valued but relatively certain gain.

Both experiments further suggest that participants' perceptuo-motor behaviour may deviate from optimality in ways not captured by Trommershäuser et al.'s (2003a, 2003b) model. Firstly, participants appeared to favour speed over precision, producing movements to near targets that were faster than necessary. Given the speed-accuracy trade off (Fitts, 1954; Schmidt, et al. 1979), such movements should decrease precision and therefore participants' ability to hit targets. The model fails to capture such apparent satisficing as it assumes that people move as to maximize precision. Secondly, participants appear to relax their precision criteria when aiming for larger targets. As Trommershäuser et al.'s model assumes that precision is maximized the model does not penalize participants for this. If participants do not always minimize movement error in perceptuo-motor tasks, an optimal model assuming that they do may make them appear more optimal than they actually are.

2.5 The effect of task parameters on performance metrics

The key result of Experiment 1 and 2 was that seemingly innocuous changes in task parameters, such as target size, can result in very different views on optimality. Next we simulate the effects of changes in task parameters, separately for the implicit and the explicit choice component, to explore in greater detail how such changes might affect participants who deviate from optimality.

2.5.1 Task parameters and the optimality of aim point choices

Fig. 2.8 illustrates the effect of changing target size on the implicit choice component. Panel A shows the optimal aim point (cross) with sample hit points (grey discs) as well as two sub-optimal aim points (triangle and square). As target size increases, naturally so does the likelihood of hitting the target (Panel B), whether you are optimal (full line) or sub-optimal (triangles & squares). Panel C shows the hit probability for the two sub-optimal aim points as proportion of the optimal hit

probability (i.e., as efficiency). It appears that sub-optimal aiming becomes less costly in terms of absolute efficiency as target size increases.




Fig 2.8. Effects of changing task parameters on implicit choice. A) Target and optimal and sub-optimal aim points (with a hypothetical response distribution [grey discs]). B) Hit probabilities for each of the three aim points: optimal (full line), small deviation (triangles), and a large deviation (squares). C) Efficiencies (hit probabilities normalized by optimal hit probabilities) for the two sub-optimal aim points in Panel A and B. D) As Panel C) but now with the lower 95% CI of optimal performance.

However, as noted in the description of Trommershäuser et al.'s paradigm, whether or not behaviour is considered optimal depends not on absolute efficiency, but the relationship between absolute efficiency and the variability of the optimal agent.

Panel D) shows the lower 95% confidence bound on the optimal agent's hit efficiency (dashed line). When either of the two sub-optimal aim points (triangles, squares) results in efficiencies above the dashed line, participants would be classed as optimal. Conversely, efficiencies lower than the dashed line implies that participants are sub-optimal. As can be seen (Panel D), smaller targets result in more variable optimal agents (wider CI's). This means that small targets allow for greater deviation from the optimal aim point before participants are classed as sub-optimal.

How do these simulations fit with the results of Experiment 1 and 2? Targets in Experiment 1 were smaller than targets in Experiment 2. This means that sub-optimal participants should have been more likely to be classed as optimal in Experiment 1. This is precisely the pattern of results obtained. There were significantly more optimal participants in Experiment 1 than in Experiment 2, and this difference appeared driven by differences in confidence intervals rather than differences in absolute efficiencies. We return to the issue of the confidence interval difference between Experiment 1 and 2 below, and show that changes in target size is likely to have accounted only for a small part of the total effect.

2.5.2 Task parameters and the optimality of target choices

Fig. 2.9 illustrates the effect of changing hit probabilities for the other component of the task: target choice (the explicit component). Fixing the large target's distance and size, we increase the size of the smaller target. The simulated large target is sufficiently large so to achieve a hit probability of near 1. Panel A illustrates the effect of this manipulation on hit probabilities for the small target (dashed line) relative to the large target (full line). As we increase the small target's size (increase target size ratio), it becomes increasingly easy to hit (hit probability increases).

Of course, for choosing between the small (dashed line) and the large target (full line), knowing hit probabilities is not sufficient; we need to know the rewards associated with each target. Panel B shows the number of points we can expect to earn for the respective target for the reward structure employed in Experiment 1. For a small to large target size ratio of up to $\sim .4$, the large target should be chosen (its expected value is higher). With further increases in the small target size, one should switch and choose the small target. ΔEV is the difference in expected value between the small and large target. If it is positive, the smaller target is worth more and should be chosen (if negative the large target is worth more).




Fig 2.9. Effects of changing task parameters on explicit choice. A) The effect of the target size ratio on hit probability for the small (dashed line) and large (full line) respectively. B) Expected value of the small (dashed line) and large (full line) target as a function of target size ratio. ΔEV is the difference in expected value between the small and the large target (see text for explanation). The cross represents a hypothetical choice situation in which the small target should be chosen. C) Choice predictions (as proportion small target choices) for an optimal agent (black step-function) and a less-than-perfectly sensitive sub-optimal agent (grey function).

The black step-function in Panel C (Fig. 2.9) illustrates the behaviour of an optimal participant who maximizes expected value (as in Trommershäuser et al.’s 2003a, 2003b model) and its shape is illustrative of the all-or-none prediction of maximization theories in general (e.g., expected utility theory). As can be seen, if the small target is worth more (positive ΔEV) it is always chosen (the proportion of small target choices is 1), and conversely when the large target is worth more (negative ΔEV) it is always chosen.

Although the black line illustrates normative responses, which are perfectly sensitive to ΔEV , people are unlikely to exhibit such sensitivity. Consequently, one might expect better choices when ΔEV is large (because it should be more readily apparent which of the two targets is better, see e.g., Mosteller & Nogee, 1951; see Brandstätter, Gigerenzer & Hertwig, 2008 for this idea applied to model evaluation).

The grey function in Panel C illustrates a participant who is less-than-perfectly sensitive to differences in expected value (ΔEV). The cross in Panel B and C, illustrates a particular choice situation, in which the optimal response is to choose the small target. A partially sensitive participant (grey line) will only pick the optimal target ~80% of the time – leading to a loss relative to the ideal agent (black line). From the grey function, it

should also be clear that as the absolute ΔEV becomes larger, the optimal agent and the sub-optimal agent become increasingly similar.

How does the above relate to the explicit choices in Experiment 1 and 2? In Experiment 1 the difference between the small and large target reward was greater and the target-distance differences were greater than those in Experiment 2. This should have made expected value differences in Experiment 1 larger. The mean absolute expected value difference was indeed greater in Experiment 1 than in Experiment 2 ($t(13) = 2.74$, $p = .017$, mean difference = 4.23, one outlier ~ 2.5 inter-quartile ranges from the median in Experiment 2 excluded). Thus, for someone who is only partially sensitive to ΔEV differences, Experiment 1 should be easier than Experiment 2. One indication that this was the case is the fact that biases were less severe in Experiment 1 (see Fig. 2.6).

2.5.3 Task parameters and bounds of optimal performance for whole experiments

The confidence interval on the optimal agents' performance is crucial. It is used to infer whether or not participants are optimal. Indeed, across Experiment 1 and 2, participants' absolute efficiencies were approximately equal. Yet, participants in Experiment 1 were classed as optimal, and those in Experiment 2 were classed as sub-optimal. Experiment 2 resulted in a more lenient standard of optimality as the confidence intervals of optimal agents' earnings were wider.

What accounts for the wider confidence intervals in Experiment 2? As outlined above, target size and the reward for the large target differed across experiments. A third factor is sample size. Increasing the number of unique choice options as was done here (Experiment 1 = 8, Experiment 2 = 18), whilst keeping the number of choices for each choice option constant, results in a different number of total trials. The total number of trials was substantially greater in Experiment 2 ($42 \times 18 = 756$) than in Experiment 1 ($42 \times 8 = 336$). Everything else being equal, a greater sample size leads to tighter confidence intervals. Thus, the difference in confidence intervals could potentially be accounted for by changes in target size, changes in rewards and/or changes in total sample size.

We explored the effect of these three factors by simulation. We simulated the ideal agents of Experiment 2 (tight confidence intervals) under conditions which were made increasingly similar to those of Experiment 1. To make the two maximally comparable, we selected the target locations in Experiment 2 that were most similar to

those in Experiment 1 (near and far). This also has the beneficial effect of equating the total number of trials for the simulated experiments to that of Experiment 1.

Fig. 2.10 shows the tight confidence interval reported for Experiment 2 (grey discs, ‘Exp 2’, identical to Fig. 2.7, bottom panel). The other symbols illustrate the effects of different reward and target size combinations when the number of target locations is equal (‘N’) to that in Experiment 1. For example, ‘Exp 1: N & size’, means that the total number of trials and the target size were the same as in Experiment 1. The shaded region represents the 95% confidence bound on the average lower bound on efficiency reported for Experiment 1 (Fig. 2.7, top panel). If Experiment 1 and 2 were identical, one would expect the average confidence interval for Experiment 2 to lie in this shaded region.




Fig. 2.10. The lower confidence bound of optimal efficiency as a function of sample size, target size and reward structure. The five different symbols represent the lower 2.5% bound of optimal performance as Experiment 2 is made increasingly similar to Experiment 1. The shaded area between the dashed lines represents the 95% confidence

bound (bootstrapped) on the *average* lower 2.5% bound of optimal performance in Experiment 1.

The difference between the original fit (grey discs, ‘Exp 2’) and the other fits is a measure of the effect size of a particular change. For example, it is apparent that changing only the number of trials (compare the ‘Exp 2’ to crosses, ‘Exp 1: N’), has a relatively small effect. Likewise changing the number of total trials and size (stars, ‘Exp 1: N & size’) has a relative minor effect. Making the difference in small and large rewards larger (triangles, ‘Exp 1: N & reward’), as in Experiment 1, or changing both the rewards and the target sizes (squares, ‘Exp 1: N, size & reward’), has substantially greater impact.

In fact, once the agents of Experiment 2 experience similar conditions to those of Experiment 1 (squares, ‘Exp 1: N, size & reward’), 5 of 8 simulated bounds lie within the confidence interval of the average *actual* bound of Experiment 1. The slight underestimation of variability relative to Experiment 1 (shaded area) is likely due to the fact that far targets were nearer in Experiment 2 (this difference could not be simulated). Because targets were nearer, they also were easier to hit (a greater proportion had hit probabilities close to 1, and fewer close to .5), and therefore resulted in less variable gains, which lead to tighter confidence intervals.

It is perhaps surprising that the effect of doubling the number of trials has such a relatively minor effect on the width of the confidence intervals. However, for an entire experiment, what matters is the variability on the gain achieved in the whole experiment. This variability depends not only on the confidence intervals for hit-probabilities of particular targets (as illustrated in Fig. 2.8), but also on the specific combinations of rewards, penalties and hit probabilities across targets (as shown in Fig. 2.10).

To illustrate, consider a task in which an optimal participant attempts to hit two different targets. In one experiment, the rewards for hitting the targets are 259 and 10. In another experiment the rewards for the targets are 115 and 100. In both experiments the penalty for missing is -5, and the higher valued target is harder to hit than the lower valued target (hit probability of .5 vs. .8).

An optimal participant is expected to earn the same number of points in both experiments (134). However, the confidence interval on the expected gain will be very different. In fact, the former scenario (259 vs. 10) will result in confidence intervals almost twice the width (~77 vs. ~44) of the latter scenario (115 vs. 100). This example

illustrates that although one can, as was shown above, illustrate some of the potential problems of categorizing participants as optimal and sub-optimal by breaking down the effects of particular changes in task parameters; the final verdict on whether people are optimal or not, depends on task parameters that interact in ways not easily captured by such modelling.

2.6 General discussion

2.6.1 Summary of empirical results

It has been suggested that the perceptuo-motor system makes optimal decisions in tasks that require both explicit target choice and implicit aim point choice (Trommershäuser et al., 2006; Seydell et al., 2008). Using a novel perceptuo-motor decision task, we found that this was the case for one particular set of task parameters (Experiment 1), but not for another particular set of task parameters (Experiment 2). Even in Experiment 1, where participants were mostly optimal, participants' efficiencies were consistently lower than lower than 1. That is, efficiencies did not cluster around 1 as expected if participants had been optimal.

We argued that the likely origin of the difference between Experiment 1 and 2 was a more lax criterion (wider confidence intervals) of optimality in Experiment 1. The more lenient criterion appeared to have been caused by seemingly innocuous changes in task parameters, such as changes in target size and target rewards - changes which do not affect the expected performance of an optimal agent – yet had dramatic effects on whether or not participants were classed as optimal. Through explorations of the optimal model we showed that changes in parameters across Experiment 1 and 2 were such that Experiment 1 is likely to be easier for (sub-optimal) participants than Experiment 2.

2.6.2 Apparent relaxation of precision criteria when aiming for large targets

In addition to task-parameter dependent optimality, our results suggest that people are sub-optimal in ways not captured by the implemented model (Trommershäuser et al., 2003ab). Participants reached with greater precision to small targets than to large targets, which suggests that humans sometimes satisfice rather than maximize precision. In a small control study (see Supplementary Materials: Chapter 2 for details), we tested whether participants could reach with equal precision to small and large targets when they were explicitly asked to do so. Under these conditions, three of five tested participants reached with equal precision to small and large targets. Given this

data, the null hypothesis of equal precision was more than three times as likely as the alternative hypothesis that the precision was unequal (JZS Bayes Factors > 3), with the evidence for the two other participants being inconclusive. Consequently, the failure to reach to small and large targets with equal precision in Experiment 1 and 2 does not appear to be due to a capacity limitation. One can also show that the apparent precision-satisficing in Experiment 1 and 2 was consequential by simulating optimal agents who aim with equal precision to both target sizes. Had participants been compared to such agents, their efficiencies would have dropped significantly relative to the standard analyses presented above ($t(17) = t(15) = -3.74$, $p = .002$, mean difference = $-.02$).

Given that ours appears to be the only perceptuo-motor decision-making experiment to have looked for target-size effects, the implications of the apparent precision satisficing remains unclear. The problem participants faced in our task is different to that in most previous studies⁹. This may mean that the precision-satisficing finding is of limited applicability to earlier studies (e.g., Trommershäuser et al., 2003a, 2003b). Nevertheless, in the absence of independent assessment of the precision maximization assumption, it is at least possible that it also applies to earlier studies.

2.6.3 Apparent under-utilisation of response time

Participants evidenced another type of apparent sub-optimality not captured in Trommershäuser et al.'s (2003a, 2003b) model. As Gepshtein et al. (2007), we found that participants did not make use of all the available response time when pointing to near targets. Can this under-utilisation be explained by reference to movement duration minimisation? Tanaka, Krakauer and Qian (2006) have proposed that the motor-system minimizes movement duration whilst keeping movement variability below a criterion determined by the task at hand. For example, putting a key in a lock seemingly requires a high precision, whereas picking up a sock off the floor does not. If precision need not be maximized, the latter task allows faster movements. Their model could potentially explain both the failure to maximize precision and the failure to maximize response time.

⁹ In our task, the control problem is essentially three-dimensional, whereas Trommershäuser et al.'s task is essentially two-dimensional. Pointing movements in their tasks are away from the body towards a screen facing the participant. This implies that controlling variance in the depth plane is relatively unimportant (providing the trajectory is relatively perpendicular to the screen as the finger approaches it) as movements are stopped by the screen. Our tasks involved pointing along a, in the depth plane, curved trajectory (from a right-side dock, towards the body, away from the body, to a left-side target). This meant that although the targets were 2D (as in Trommershäuser et al.'s studies) – variance in the depth plane mattered (poor depth control should result in over or undershooting).

However, it appears that at least some participants traded-off reaction time and movement time. These participants initiated reaches to harder-to-hit targets (e.g., far targets) faster than they initiated reaches to easier-to-hit targets (see see Supplementary Materials: Chapter 2 for details for data suggestive of this). This is interesting as it suggests that although response time usage is not maximized, the perceptuo-motor system is somewhat compensating for the extra movement time needed for targets that are hard to hit (but there were signs that the reaction times to smaller targets were actually slower than to larger targets – questioning the generality of this point).

Our results suggest that neither reaction time nor movement time is minimized. If the system was designed to move as fast as possible given a set criterion (cf., Tanaka et al.), it would presumably initiate movements as fast as possible as this would allow for faster motor action (all other things being equal). The fact that reaction time was modulated by target distance means that it does not do this. Instead, trading off movement time with reaction time, and satisficing precision, may be a result of optimizing a more complex function (perhaps one that can be adjusted on the basis of task demands as suggested by Todorov and Jordan's [2002] optimal feedback control theory).

2.7 Theoretical implications

The main implication of these experiments, as we will argue, is that the standard method of analysis (e.g., Trommershäuser et al., 2003a, 2003b; Wu et al., 2006) where participant's performance is compared to an optimal agent, based on characteristics of that participant's performance, does not necessarily result in an absolute performance standard. Consequently, such analyses do not seem to support conclusions about optimality "in general". Instead, statements about optimality are specific and conditional: behaviour is optimal *given* a task of this difficulty, and *given* these capacity constraints included in the optimal agent. In this case, however, it is unclear how such analyses could support comparisons across tasks (whether within or between cognitive domains). In other words, it may not be appropriate to use this method to make unconditional inferences about the optimality of behaviour.

2.7.1 The Problem of Task Difficulty

We have shown that it is possible to make people appear optimal, or sub-optimal, by seemingly innocuous changes in task parameters. Thus, it seems that models such as the one used here are sensitive to task difficulty. When a task is "easy" performance is

good, when a task is “hard” performance is bad. If one wants to make claims such as “system X is optimal” this is unfortunate. It would seem a desirable property (of optimal models) that the classification of a studied system does not change as a function of minor task changes. Secondly, even if this were to be viewed as a non-problem – the problem of specifying a “standard” task against which performance confidently can be assessed remains. Which target size, for example, represents a suitable difficulty level when assessing aim point choice efficiency? Unless task difficulty can be defined independently of tasks, comparing performance within, or between, modalities seems difficult. Naturally, if whatever makes the task harder for participants were to be correctly modelled, performance would no longer be a function of task difficulty, but, of course, this does not make the ideal standard absolute. Instead, it highlights the conditional nature of such models.

2.7.2 On the use of optimal agents to infer unconditional optimality.

At first glance, quantitative ideal observer methods might seem suited to take into account task difficulty. In fact, ideal observer methods have been promoted on the grounds that they provide an absolute standard (e.g., Shimozaki, Kingstone, Olk, Stowe & Eckstein, 2006). Note, however, that this standard is critically dependent on the constraints imposed on the ideal agent. An ideal agent, whose constraints *exactly* match those of the participant, is likely to match the participant’s performance exactly. Likewise, an ideal agent who incorporates none of the constraints is likely to be vastly superior to the modelled system. Performance is not absolute, but relative to the constraints built into the model. Importantly, if built-in assumptions are unjustified, they may cause overestimation, or underestimation, of performance independently of task difficulty.

This is in no way to say that ideal observer analysis is flawed as a method. Ideal observer analysis is an extremely useful analytical tool, as evidenced by the arguably tremendous success it has had. It can, for example, be used to chart the efficiency of a system (e.g., Barlow, 1962) or it can be used to constrain the search for plausible models (Schrater & Kersten, 2002).

2.7.3 General implications for inferences about the optimality of behaviour

The above points arguably apply with equal force to the cognitive domain. One may argue that because one can show fundamental deviations from normative standards under simple experimental conditions (e.g., choices between two options), the problem

of task difficulty does not arise. To drive the point home, one may further argue that many real world problems are orders of magnitude more complex. However, this argument fails to recognize that many of the experimental problems in this literature differ very little in terms of expected outcome. If experiments, in addition, have low ecological validity they may not tap processes that are most likely to have been optimized (either through learning or evolution).

To illustrate, consider a choice pattern that violates maximization of expected value¹⁰. When asked to choose between a gamble that yields \$2500 with a .33 probability, \$2400 with a .66 probability and \$0 with a .01 probability and a gamble that yields \$2400 with certainty, most tend to pick the latter (Kahneman & Tversky, 1979, pp. 265–266). The expected value of the former is \$2409 and the expected value of the latter is \$2400. Note that the expected loss of choosing the modal response is only 0.4%. Presumably, people would not select the lottery with the lower expected value if the difference were much larger. Furthermore, assuming noisy computational processes (Faisal, Selen & Wolpert, 2008), people might not even be able to distinguish between expected values that differ very little. Thus, it may be argued, that the decision is a hard one.

In general, any decision problem could be made difficult enough to be unresolvable by an actual physical system. Trivially, for example, differences in expected utility might be made so slight (for example, present in decimal places only) as to exceed the resolution of the system. Typically, however, such limitations will not be perceived as interesting. Rather, limitations are of note typically only where we have reason to believe that the system *could*, or *should*, be able to deal with the problem at hand. This is typically not articulated explicitly, but permeates all research on human rationality. Experimental demonstrations of norm violations such as base rate neglect, the conjunction fallacy, framing effects, or logical reasoning errors generate widespread interest because they are perceived as ‘gross’ errors. Subsequent research then typically seeks to at least partially restore the case for human rationality by demonstrating that unnatural or misleading problem formulations are to blame for poor performance, or that the error in question, ‘in the real-world’, is ultimately not a costly one (e.g., Hilton, 1995). In other words, subsequent research challenges the perception that participants *should* readily be able to avoid these errors.

¹⁰ This exposition relies on expected value, the original problem deals with expected utility. However, as the perceptuo-motor decision-making literature has relied on expected value (gain) and not utility, expected value is used as an example here.

Both demonstrations of rationality and of irrationality, it seems, are inherently set against background expectations that make the observation interesting, whether as a surprising failure, where success seems reasonable, or as striking performance where a task seems difficult. It would seem that there are no absolute standards of rationality in practice. This is in no way to say that research on the rationality or optimality of human performance is not informative. However, the limitations of the kinds of statements that are being made must be considered.

2.8 Summary

We have argued that statements to the optimality of a given system can be made only conditionally (at least when standard analyses are used). Specifically, task difficulty and modelled constraints influence the classification of a system as optimal or sub-optimal. In order to make performance comparisons within or between systems, one needs an absolute standard of performance. Alternatively, one can try to equate task difficulty and modelled constraints across tasks as far as possible. The latter solution solves the problem of conditional optimality by attempting to equate the conditionality of task performance. Whilst this does not seem to allow unconditional statements such as “system X is optimal”; it does allow comparisons across modalities to the extent that they can be presented with “the same” problem. Wu et al. (2009) represents an interesting example of this latter approach. Interestingly, they found little evidence for previously made comparative claims favouring the performance of perceptuo-motor decision-making over higher-level cognitive decision-making¹¹.

In conclusion, presently, there seems to be little basis for the claim that human perceptuo-motor decision-making is optimal and that human cognitive decision-making is not.

¹¹ Wu et al. (2009) conducted two experiments. In the first they found no evidence that the degree to which the independence axiom is violated differs between perceptuo-motor choices and classical paper-and-pencil type decisions. In a second experiment they found that, for a particular parameterization of cumulative prospect theory, the probability weighting function parameters differed between the two types of decisions. Whilst this is interesting in of itself, the mere fact that probabilities need to be weighted in the first instance means that decisions were sub-optimal (albeit in different ways). See Wu et al., 2009 for details.

3. Interlude 1

In the previous chapter, we showed that one can make the perceptuo-motor system appear optimal, or sub-optimal, by making seemingly innocuous changes to a perceptuo-motor decision-making task. For example, by making the possible distance differences between targets larger, we increased the difference in expected value between pairs of targets. This change does not affect the performance of an agent who behaves according to the optimal model of Trommershäuser et al. (2003a, 2003b). However, as we argued, one might expect this change to affect human participants. The result that human performance seemingly was affected by such changes, yet the model not, suggests that Trommershäuser et al.'s optimal model may need extending to fully capture human behaviour.

Some changes across tasks did, however, affect the behaviour of the optimal agent. Importantly, these changes caused the earnings of the optimal agent to become more variable. The standard method of assessing performance is to compare participants' actual performance to the distribution of the optimal agent's earnings (see e.g., Trommershäuser et al., 2003a Wu et al., 2006). If participants' earnings lie outside the lower 2.5 percentile of this distribution they are classed as sub-optimal. Therefore, if one can manipulate the standard deviation of that distribution, without otherwise affecting participants' absolute efficiency (i.e., their earnings relative to the expected earnings of the optimal agent), it should be possible to make people look either optimal or sub-optimal, despite the fact that their actual performance relative to the ideal observer has not changed. The widening of the confidence interval in Experiment 2, relative to Experiment 1, without a noticeable difference in *absolute* efficiency levels for our participants, suggests that we succeeded in doing just that.

These results were worrying to us. The fact that one can change the classification of the perceptuo-motor system from optimal to sub-optimal, by making such small and apparently harmless changes to tasks, implied difficulties for the assessment of the perception-cognition gap. For how can one justifiably compare performance across perceptual and cognitive tasks, if the standard of optimal performance is relative, not absolute - *even* within a single domain?

The next chapter represents an attempt at getting around the problem of relative performance standards. The central idea was to use a decision task that might be viewed as modality independent. One such task is making decisions about how much time to spend. To decide how much time to spend on a given task wisely, you need to know how costly it is to get the task wrong, how rewarding it is to get it right, and how your

task performance changes as a function of how much time you spend on the task. Importantly, deciding how much time to spend on a given task is a decision problem that arises both for tasks that are primarily perceptual and for tasks that are primarily cognitive.

In the next series of experiments we used decisions about time to first explore perceptual performance in some detail. We then contrasted perceptual and cognitive performance. That is, we compared the two in a situation in which the decision task itself (time decisions) and the reward information (points converted to money) was *identical* across “cognition” and “perception”. What identified the tasks as either perceptual or cognitive was how the underlying knowledge required to perform the decisions was derived (from cognitive or perceptual experience). In other words, the probabilities (here in the form of accuracies) came from either the perceptual or the cognitive domain, but the value information was identical across domains as was the type of decision. The following experiments might therefore be viewed as a test of whether the format in which uncertainty information is presented affects decision performance.

4. Knowing When to Move On: Cognitive and Perceptual Decisions in Time¹²

A hunter-gatherer is tracking prey through a forest. In his path lies a pond. Does he think that it would take him longer to pass it by going to the left or to the right? The time it would take him to answer that question with certainty is likely longer than the time lost by making the wrong choice. Thus, he deliberates only briefly before going left. Then, coming upon dense undergrowth he sees movement. It is an area known for dangerous predators. So, before proceeding, he spends a considerable time making sure that he knows what caused the movement. These decisions are about how much time to spend on the task at hand. Good performance depends on, among other things, taking the cost of errors into account and having knowledge of one's own task performance.

Given that decisions are an integral part of life, it seems plausible that humans have acquired, through a combination of learning and evolution, excellent decision making skills. In line with this idea, human perceptuo-motor and perceptual decisions do often appear near-optimal (Trommershäuser, Maloney & Landy, 2003a, 2003b; Whitely & Sahani, 2008; Navalpakkam, Koch & Perona, 2009; Navalpakkam, Koch Rangel, Perona, 2010).

We tested whether people also can be near-optimal when they make decisions like those faced by the hunter-gatherer in the scenario above. First we explored people's ability to make timing decisions when the underlying task is perceptual. That is, when the crucial knowledge, of how performance changes as a function of time, is derived from perceptual experience. We found that performance was very good. Remarkably, our participants achieved this level of performance in the absence of any feedback.

Although perceptual and perceptuo-motor decisions often appear near optimal, extensive research exemplified by the work of Kahneman and Tversky (Kahneman & Tversky, 1979; Kahneman, 2003), suggests that human higher-level decision making is far from optimal. The apparent performance dissociation between high- and low-level decisions has not gone unnoticed (Trommershäuser, Landy & Maloney, 2006; Trommershäuser, Maloney & Landy, 2008).

Intrigued by this perception-cognition gap, we wondered whether it might apply to timing decisions. Specifically, performance might be affected by the “modality” of the underlying task. Decisions for which the required knowledge is derived through perceptual experience might be better than those for which the knowledge is derived

¹² A version of this chapter has been accepted for publication in *Psychological Science*.

from cognitive experience. We were not able to confirm such a difference. In fact, timing decisions were near-optimal whether the underlying task was of a perceptual or more cognitive nature. The highly similar performance suggests that knowledge can be acquired, and used for timing decisions, in an equally proficient way whether it is derived through perceptual or cognitive experience.

4.1 Experimental paradigm

Our experiments involved two stages. The purpose of the first stage was to estimate how accuracy changes as a function of time spent on the given task. The tasks we employed were like the initial ‘shortest path around the pond’ example. There were two possible outcomes: success or failure in picking the right answer. If only a very brief moment is spent on such a task, performance will be very poor (accuracy $\sim .5$). As the time spent on the task increases, so should performance - until performance reaches a plateau where further increases in time do not yield improvements (e.g., when accuracy = 1). We assessed participants’ accuracy for six different task durations (Fig. 4.1A, dots). The time available was constrained by forced deadlines (Schouten & Bekker, 1967). A first tone was presented at stimuli onset, another tone was presented half way to the deadline, and a final tone was presented at deadline. By manipulating deadlines we were able to sample times that resulted in a range of accuracies (from $\sim 50\%$ correct to asymptote). We used this data to estimate the relationship between time and accuracy by fitting a function (Weibull, 1951, Fig. 4.1A, line, see “SI Methods: Fitting the Weibull function to accuracy data” for details). This function describes the time-accuracy relationship well (see “SI Methods: Assessing goodness of fit”).

In the decision making stage, participants were given an overall time interval within which to complete as many individual trials of the task as they saw fit (Fig. 4.1B, top). Participants were told the values of correct and incorrect responses in points, were paid a performance related bonus and were instructed to earn as many points as possible. To maximize their earnings they had to choose, among the many possible average response times, the one that maximizes reward (Fig. 4.1B, bottom, see “SI Methods: Mathematical formulation of the decision problem” for a formal description). To ensure that we were studying on-line decisions (rather than, for example, learning of decision criteria or stimulus-response contingencies) no feedback was provided.¹³

¹³ The absence of feedback here is worth stressing, because the many studies of timing decisions (Bogacz, 2007), largely focussed on testing specific models (Grice, 1968; Ratcliff, 1978) and not the optimality of timing decisions in general (but see Bogacz, Hu, Holmes & Cohen, 2010), have involved extensive feedback (as have perceptuo-motor studies with time-based cost functions, see Trommershäuser et al.,

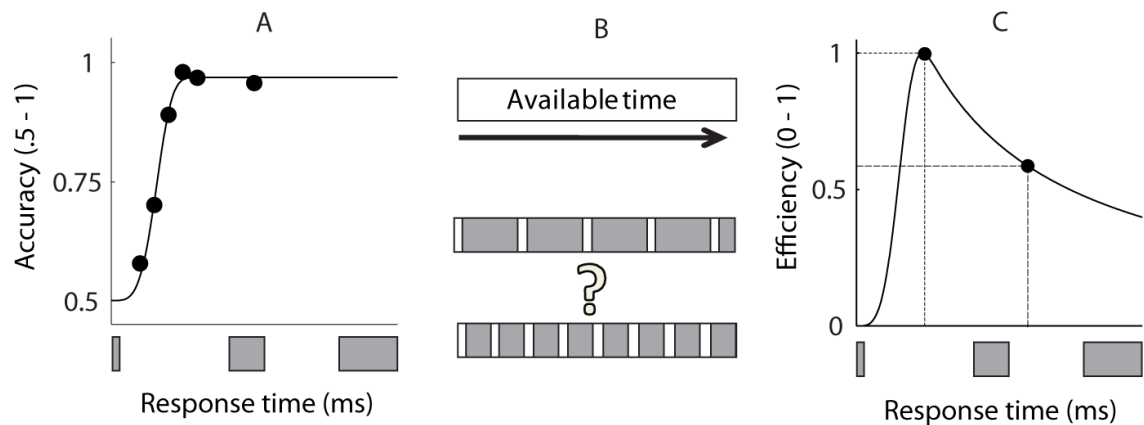


Fig. 4.1. Experimental paradigm. (A) Data from the assessment stage fit with a Weibull function. (B) Decision making stage (top), choice of average response time (bottom). (C) An efficiency function with two hypothetical timing choices.

To evaluate participants' choices we mapped average response times onto efficiency functions (Fig. 4.1C). An efficiency function describes how well one would be doing (on average) for a given response time. An efficiency of 1 means that one is earning as many points as one possibly could (and money, as bonus pay was proportional to efficiency). Responding at a faster, or slower, time than optimal, will result in a lower efficiency (and less money earned).

Efficiency functions were derived by taking into account a participant's time-accuracy function (Fig. 4.1A) and the rewards/penalties along with the refractory period between decisions. By mapping response times onto independently assessed efficiency functions (see "SI Methods: Extrapolation from Weibull fits to choice data" for validation) decision making performance can be assessed across a range of underlying tasks. Our approach of mapping response times onto empirically assessed functions is essentially model-free.

4.2 Timing decisions based on perceptual information

A basic question is whether people's choices relate to the efficiency functions in any meaningful way. We therefore began with a systematic exploration of decision

2008 for an overview). Feedback is known to alter patterns of decision (e.g., Barron & Erev, 2003; Camilleri & Newell, 2011) and has been found to improve performance in the supposedly weaker 'cognitive' domain (see e.g., Chu & Chu, 1990; Shanks, Tunney & McCarthy 2002; Jessup, Bishara & Busemeyer, 2008). Hence authors reporting optimality in the perceptual domain have sought to rule out learning from the extensive feedback provided as an explanation (e.g., Trommershäuser, et al, 2003a; Whitely & Sahani, 2008, but see e.g., Navalpakkam et al., 2009), setting these studies apart from decision from experience studies that have evaluated the effect of perceptual uncertainty on decisions but not sought to rule out learning-based explanations (Shafir, Reich, Tsur, Erev & Lotem, 2008).

making performance in the context of a low-level visual task - motion discrimination (e.g., Shadlen & Newsome, 2001, see “SI Methods: Additional details”). In this task, many dots move in random directions across a screen. A proportion of these dots, however, move coherently in one of two possible directions. The task is to judge the direction of the coherently moving dots. Intuitively, this task might be likened to the task of judging the direction of the wind by looking at how raindrops fall.

Using this task we manipulated two of the factors that influence the shape of efficiency functions: task difficulty and relative rewards. If people’s behaviour is well described by the optimal model, one would expect them to be sensitive to such manipulations. For example, if a particular manipulation shifts the peak to the right (slower response required), one would expect people to also shift to the right (slow down). Importantly, if people’s behaviour is well described by the optimal model, one would expect people’s timing choices to be near the optimal ones.

4.2.1 Experiment 1: Motion discrimination & task difficulty

A change in task difficulty will produce a change in the accuracy function (easy=full line, hard=dashed line, Fig. 4.2A), which in turn influences the shape of the efficiency function (Fig. 4.2B). An increase in task difficulty generally shifts the peak of the efficiency function to the right, so that a slower response time is required (Fig. 4.2B). Using a neutral reward structure (reward = 1, penalty = -1), we manipulated task difficulty by changing the proportion of coherently moving dots (“easy” = 70%, “hard” = 20%). Are people sensitive to changes in task difficulty, and if so do they choose response times that coincide with the peak of the efficiency functions?

4.2.1.1 Methods. Six members of the School of Psychology’s participant panel took part in one learning session (30 min) and two experimental sessions (30 min each) in exchange for £6/hr and an additional performance related bonus (average achieved efficiency * £6). In the learning session participants learned what buttons to press by doing the motion discrimination task with no time limit, no rewards or penalties, and auditory cues for correct and incorrect responses. They also performed the assessment stage of the experimental session in order to practise the timing requirements.

The experimental session involved two stages: assessment and decision making. The assessment stage (see “Experimental paradigm” above) involved the assessment of task accuracy as a function of time spent on the task (with deadlines manipulated across blocks of trials). The decision-making stage involved 2 two minute periods (spread across two sessions) for each task difficulty. Participants were informed that they would

earn 1 point for each correct response, -1 point for each incorrect response, that the goal was to earn as many points as possible and that the more points they earned the more extra money they would receive. Importantly, there was no feedback on whether a given response was correct or incorrect, or on points accrued.

For data analysis, average response times were computed from raw data. Efficiency was defined as re-scaled (confined to lie between 0 and 1) expected value. Ninety-five percent confidence intervals for mean response times and efficiencies were computed using bootstrap methods (Efron & Tibshirani, 1993).




Fig. 4.2. Blocked task difficulty experiment (dashed lines = hard task, full lines = easy task). (A) Participant 2’s accuracy functions. 15 of 18 easy-hard accuracy function pairs were statistically different (includes comparisons for Fig. 4.5, see “SI Methods: Was the hard-easy manipulation successful?” for details). (B) Participant 2’s efficiency curves with choices (circles). (C) Efficiency for each task difficulty for Participant 1-6. Gray bars = hard task. White bars = easy task. Error bars are bootstrapped 95% CI’s. The dashed line shows an efficiency level of .95.

4.2.1.2 Results & discussion. Fig. 4.2C shows efficiency as a function of task difficulty for each participant. As can be seen, most participants responded appropriately to the difficulty manipulation. In fact, four out of six participants have efficiencies near 1—regardless of whether the task is hard (gray bars) or easy (white bars). From the individual efficiency functions (Fig. S4.1), it is evident that participants made choices that indicate that they are sensitive to changes in task difficulty and are consistent (4/6 participants are near the peak for both hard and easy choices [one of these participants shows an aggressive bias] and 1 participant is consistently cautious). Thus, people appear to have good internal estimates of task performance and can use these to make efficient decisions.

4.2.2 Experiment 2: Motion discrimination & changes in rewards and penalties

Like changes in task difficulty, changes in rewards and penalties influence the shape of the efficiency function. Increasing the penalty relative to the reward generally shifts the peak to the right (slower response times required), whilst decreasing it shifts the peak to the left (faster response times required). Past studies utilising reward structures with penalties and rewards of unequal magnitude have produced contradictory findings. Some results suggest that people do respond near-optimally (e.g., Trommershäuser et al., 2003a; Navalpakkam et al., 2009) whereas others suggest that people deviate from optimality (e.g., Green & Swets, 1966; Ulehla, 1966; Maddox, 2002).

To establish that people are sensitive to changes in rewards and penalties, we initially compared choices for the neutral reward structure used above, to a standard reward-only condition (as in e.g., Shadlen & Newsome, 2001; Bogacz et al., 2010) and its reciprocal (penalty-only) in a pilot study¹⁴. The results suggested that people are sensitive to changes in rewards and penalties (e.g., if required to slow down they generally do), but that this sensitivity is somewhat limited (e.g., they slow down, but not by the right amount, see Fig. S4.2-3 for details).

To explore this partial sensitivity, we sought to determine whether people are sensitive only to rank order information about values, as has been assumed in some models of decision making (Stewart, Chater & Brown, 2006), or whether they are also sensitive to the absolute magnitude of values. To this end, we presented participants with reward structures where the reward was always 1 point, but the penalty changed across consecutive conditions in a descending order. Some participants (1-5) received a strong manipulation: -24, -18, -12 and -6. Others (6-10) performed under a weak negative manipulation: -6, -4.5, -3 and -1.5. Based on the pilot study, we expected that people would appropriately choose shorter response times as penalty levels decrease (a within-subject effect). However, for the strong vs. weak between-subject manipulation there were at least two possible outcomes. If absolute magnitude matters, as it does in the optimal model, the strong manipulation should result in longer response times than the weak manipulation (and the two -6 conditions should overlap). If, on the other hand, only the ranks of values determine peoples' choices, the two manipulations (strong and weak) should yield similar results.

4.2.2.1 Methods. Ten members of the School of Psychology's participant panel took part in one learning session and one experimental session (each 60 min). Pay was

¹⁴ Conducted in conjunction with Experiment 4 reported on below.

identical to Experiment 1. The learning and the experimental sessions were similar to those of Experiment 1, with the exception that participants now experienced one task difficulty (coherence = 25%) and four different reward structures. Participants were randomly allocated to either the strong or the weak manipulation and completed two two-minute decision periods for each penalty level (presented in a descending order).

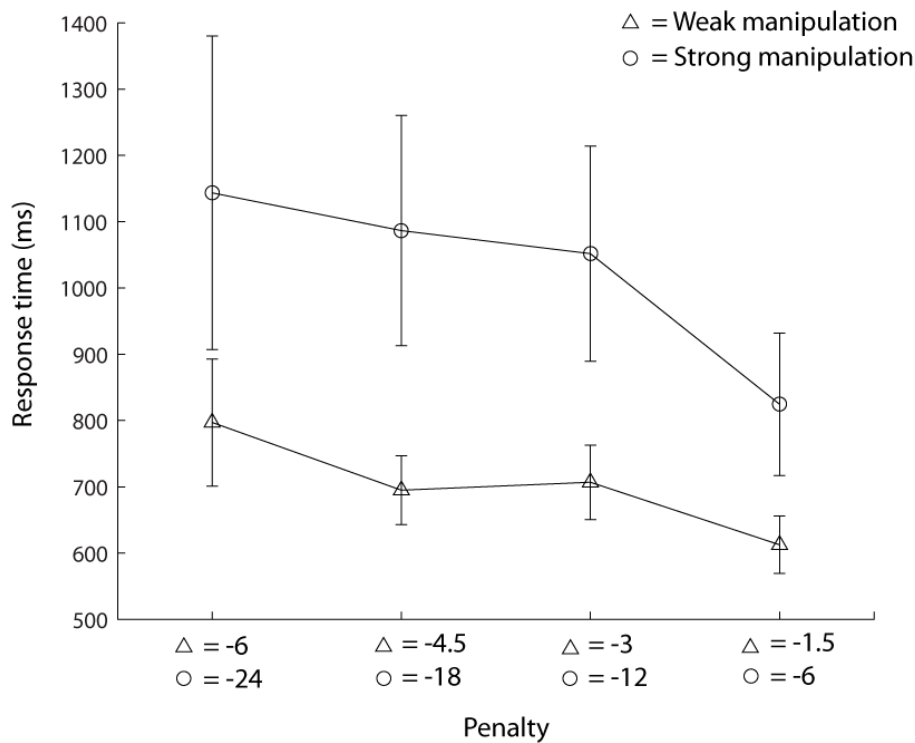


Fig. 4.3. Ordinal reward manipulation. Average response time as a function of penalty level (x-axis) and between-subject penalty manipulation (triangles = weak, circles = strong). Error bars are standard errors suitable for between-subject comparisons.

4.2.2.2 Results & discussion. First let us consider the average response times. As expected (see Fig. 4.3), participants appeared to speed up when penalties decreased. This within-subject trend was confirmed by a mixed-ANOVA ($F(3, 24) = 8.29, p = .001$, $MSE = 13426$; linear trend $F(1, 8) = 10.4, p = .012, MSE = 28301$). Comparing the average response times across the between-subject magnitude manipulation (compare triangles to circles in Fig. 4.3); it appears that higher penalties resulted in longer response times as predicted (a marginal effect; $F(1,8) = 3.36, p = .0502, MSE = 312393$, one-tailed). Furthermore, the strong -6 manipulation overlaps almost perfectly with the weak -6 manipulation - as expected if people were sensitive to absolute magnitude. The group analyses therefore indicate that people change their responding as a function of

penalty and (somewhat less conclusively) suggest that they are sensitivity to absolute magnitudes.

Let us now turn to the more important question of how participants' choices relate to the peak of the efficiency functions. As in the pilot study biases were evident. Participants slowed down, but tended to overreact to large penalties slowing down more than appropriate (see Fig. S4.4 for details). However, due to the shape of the underlying efficiency functions, these biases were not costly (see also Green, 1960; Winterfeldt & Edwards, 1968). The average efficiency for the penalty manipulation for the strong condition was .966 (SE=.018) and the average efficiency for the weak condition was .969 (SE= .011).

4.3 Timing decisions based on perceptual and cognitive information compared

Thus far, we have established that people are able to make perceptual time allocation decisions in the absence of feedback. They can make timing choices that are close to optimal whether the task is hard or easy. The ability to respond to unequal penalties and rewards appears good but biases are evident. We went on to compare timing decisions when the information they were based on was derived from low-level and higher-level tasks. In making these comparisons we restrict ourselves to a reward structure with equal rewards and penalties (1 and -1 respectively), for which choices were efficient and unbiased¹⁵.

4.3.1 Experiment 3 – Motion discrimination, mental arithmetic & mental rotation

To evaluate the effect of modality, from which the crucial information was derived, on timing decisions we employed two additional tasks (see SI Methods: Detailed methods for details). The first was a mental arithmetic task. It involved judging whether the sum of two integers was smaller or greater than 100. The second was a standard mental rotation task (Shepard & Metzler, 1967). This task involved judging whether two three-dimensional figures could be rotated mentally to bring them into alignment and incorporates both perceptual and cognitive components¹⁶.

¹⁵ See Fig. S4.5 for evidence that the bias shown under unequal penalties and rewards (Fig. S2-S4) is not restricted to the perceptual domain.

¹⁶ The distinction between high and low processes/systems/tasks is fuzzy. In the limit, all low-level tasks incorporate some high-level components (and vice versa). We did not attempt a formal distinction, but simply used tasks that are conventionally considered one or the other. Motion discrimination is a widely used psychophysical task (e.g., Shadlen & Newsome, 2001), mental arithmetic is a core cognitive ability (see e.g., Wechsler Adult Intelligence Scales), and mental rotation would seem to sit in between, as it is

4.3.1.1 Methods. Five members of the School of Psychology’s participant panel took part in two learning sessions and two experimental sessions (60 min each). Pay was identical to Experiment 1. The learning and the experimental sessions were similar to those of the earlier experiments with the exception that participants now experienced three different tasks under one task difficulty and one reward structure (reward: 1, penalty: -1). Per experimental session, participants completed one two-minute decision period for all three tasks (in random order). Experiment 3 also included a pilot designed to test the effects of relative rewards (reported in the introduction to Experiment 2 and Fig. S4.2-3, S4.5¹⁷).




Fig. 4.4. Decision making ability across three tasks. (A) Efficiency function for each participant (rows) and each task (columns). Circles represent response time choices. (B) Efficiencies across the three tasks for each participant (bars ordered as columns in A). The dotted line represents an efficiency level of .95. All error bars are bootstrapped 95% CI’s. Some error bars are too small to be visible at this scale.

4.3.1.2 Results & discussion. Fig. 4.4A shows efficiency curves for motion discrimination, mental arithmetic and mental rotation respectively (columns) for each of the five participants in Experiment 3. As can be seen, the good performance for motion discrimination shown in Experiment 1 was replicated in this experiment (Column 1). Importantly, participants’ efficiencies for mental arithmetic (Column 2) and for mental rotation (Column 3) show that these tasks can also lead to very good performance. The data further suggest that participants’ choice consistency goes beyond specific tasks

associated with cognitive ability (Ozer, 1987), but seems to engage lower-level systems also (Cohen, Kosslyn, Breiter et al., 1996).

¹⁷ This pilot study also included a manipulation of the inter-stimulus interval.

(specifically, biases are evident across tasks, Participant 3 is consistently aggressive, Participant 2 & 5 are slightly too cautious). When decision making performance is compared across the different tasks (Fig. 4.4B, bars ordered as the columns in 4.4A), performance appears largely equivalent. In other words, performance was similar, whether the required knowledge was derived through perceptual or cognitive experience.

4.3.2 Experiment 4 – Dynamic changes in task difficulty

All manipulations so far have been between blocks (i.e. the experimental variables have been fixed for a block of trials). Everyday decision tasks rarely come partitioned according to task difficulty. A more realistic experiment might involve changes occurring on a *trial-by-trial* basis. Therefore, our final question was whether people can adjust dynamically to changes in task difficulty and whether they can do so regardless of the modality of the underlying task. To answer these questions, we ran variants of the motion discrimination task and the mental arithmetic task in which task difficulty changed unpredictably from trial to trial.

4.3.2.1 Methods. Six members of the School of Psychology’s participant panel took part in one learning sessions and two experimental sessions (60 min each). Pay was identical to Experiment 1. The learning and the experimental sessions were similar to Experiment 3, with the exception that participants now experienced both the motion discrimination and the mental arithmetic task under two task difficulties (easy and hard). Task difficulty for the motion discrimination task was manipulated as in Experiment 2. An easier mental arithmetic task was created by presenting integers of 5’s instead of 1’s (e.g., 35 and 60, rather than 34 and 63). The order in which the motion discrimination and mental arithmetic tasks were completed was randomized.

4.3.2.2 Results & discussion. The manipulation of task difficulty within one and the same time-period requires a slightly different approach to the assessment of performance. For every pair of “easy” and “hard” trials, one could spend less time on one difficulty level, in order to spend more time on the other (see “SI Methods: Mathematical formulation of the decision problem” for details). Fig. 4.5A shows efficiency landscapes plots that take this into account. Choices that lie in the innermost region are at least 90% efficient. As can be seen, most participants compensate for dynamic changes in task difficulty and do so regardless of whether the underlying task is perceptual or cognitive. Participant 4 is too cautious and so falls outside the innermost region, however this behaviour is consistent across both task-types and levels of

difficulty. As in Experiment 3, the absolute efficiencies are similarly high across the two task types (Fig. 4.5B).

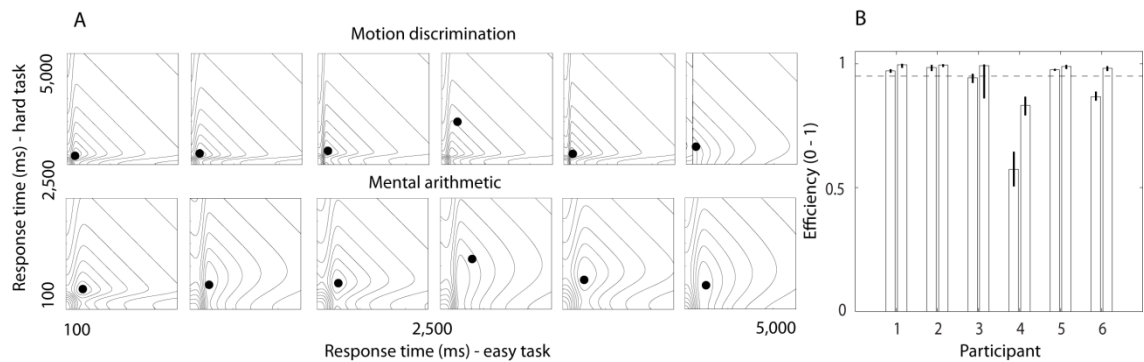


Fig. 4.5. Dynamical task difficulty manipulation. (A) Efficiency space for motion discrimination (row 1) and mental arithmetic (row 2) for Participant 1-6 (columns). Black discs are response time choices. Concentric lines represent efficiency levels. Choices within the innermost line are at least 90% efficient. Everything inside the next line is at least 80% efficient and so on. (B) Absolute efficiency levels for both types of tasks (first bar = motion discrimination, second bar = mental arithmetic). Error bars are bootstrapped 95% CI's

4.4 General discussion

We explored people's ability to make timing decisions for perceptual and cognitive tasks. Participants' choices resulted in earnings that were very close to the earnings they would have received if they had, in fact, been optimal (although, in the context of unequal penalties and rewards, participants' choices appear biased). Moreover, our participants responded appropriately to changes in task difficulty, and did so even when task difficulty changed on a trial-by-trial basis.

Decisions also appeared equally efficient regardless of whether the underlying task was perceptual or cognitive. In our experiments, the modality dependent component came from participants' knowledge about their own task performance. This knowledge was either of a perceptual or more cognitive nature. The type of decision (about time), and the reward information (abstract), was *identical* across the different underlying tasks. The highly similar performance, across lower and higher-level tasks, therefore suggests that knowledge was acquired, and used, in an equally proficient way whether it had a perceptual or a cognitive origin.

Our comparison across perception and cognition was motivated by the perception-cognition gap; low-level decisions appear near-optimal whereas higher-level decisions

appear sub-optimal. There are at least two explanations for the gap. The gap might be due to an underlying difference in competence (Trommershäuser et al. 2006; 2008). Alternatively, the gap might be a side-effect of the many methodological differences between low- and high-level decision making studies (Maloney, Trommershäuser & Landy, 2007). Differences exist both in terms of experimental methods (e.g., feedback, real stakes, Maloney et al., 2007) and in the way performance is assessed (i.e., violation of decision theoretic axioms vs. actual performance compared to optimal performance)¹⁸.

Can our results shed light on the perception-cognition gap? The answer, it seems, depends on how timing decisions are viewed. If they are viewed as products of a highly specialized system that does timing decisions only, the answer is very little (the system's ability to use perceptual and cognitive information equally efficiently notwithstanding). If, on the other hand, timing decisions are a result of a more general purpose decision system, our results would suggest that the previously reported gap, at least partly, is due to methodological differences.

Regardless of the true cause(s) of the perception-cognition gap, and the type of mechanism(s) that underlie timing decisions, our results show that timing decisions, whether based on perceptual or cognitive knowledge, can be performed near-optimally. Of course, our results do not prove that timing decisions are always near-optimal. In fact, the biases we found suggest that they are not. To return to the initial example; the hunter-gatherer, when looking for dangerous predators is likely to be too careful, but the cost of being too careful, in this situation, is much lower than the cost of being too careless. On the other hand, when choosing a way around a pond those costs are approximately equal. Judging by our results, the hunter-gatherer would have spent just the right amount of time at the pond, but a little too much time looking for dangerous predators – a sin he might easily be forgiven for.

¹⁸ Relatedly, the decision from experience literature (for a review see Hertwig & Erev, 2009) has examined methodological differences *within* the cognitive domain (e.g., Barron & Erev, 2003), finding that, for example, perceptual reward uncertainty (Shafir et al., 2008), and trial-by-trial rewards (Camilleri & Newell, 2011) affect how decisions deviate from expected utility theory. Here we have focussed on evaluating the actual efficiency with which decisions are made (as is typical in low-level studies), rather than attempting to find patterned deviations (as is typical in classical and decision from experience studies).

5. Interlude 2

The four experiments reported on in Chapter 4 showed that timing decisions can be made very efficiently. Most of our participants earned nearly as much money as they would have earned if they had, in fact, been optimal. They were sensitive both to changes in task difficulty and changes in relative rewards. When we manipulated the nature of the underlying task, participants showed near-identical performance for perceptual and cognitive tasks. This latter finding is perhaps the most interesting from the perspective of this thesis. Minimally, it suggests that decisions about how much time to spend can be made efficiently, regardless of whether the required information is derived through perceptual or cognitive experience.

The knowledge that our participants based their choices on was essentially equivalent to the probability information in standard decision-making paradigms. As discussed in *General Introduction*, there are two components to a good decision: probabilities and values. The probabilities for the timing decisions were either “cognitive” or “perceptual”, whilst the values and the decision type were identical. Of course, timing decisions require not only taking values and probabilities into account (see SI Equation 4.9), but this should make our results more and not less impressive. The previous chapter then shows that timing decisions can be performed efficiently, whether decisions involve probability information derived from low- or high-level tasks.

In terms of the overarching goal of investigating the perception-cognition gap there are two potential shortcomings in the previous chapter. A first objection is that timing decisions may be “special”. Given this possibility, one cannot, on the basis of our results, confidently infer that cognitive and perceptual decisions *in general* are equally good. Whilst this is true, it would be surprising if timing decisions were the *only* type of decisions for which the use of perceptual and cognitive information can be equally efficient. Nevertheless, intuition is not evidence.

The second potential issue is that the previous series of experiments did not involve classical cognitive decisions (as in e.g., Kahneman & Tversky, 1979) as a standard for comparison. That is, the tasks did not also include decisions for which the uncertainty information is provided in numerical format. This is a minor problem. If it were the case that decisions relying on acquired cognitive information *in general* were as good as those based on acquired perceptual information, but both were better than decisions based on numerical probabilities, this, in our minds, would be evidence against a general perception-cognition gap.

Nevertheless, the last piece of experimental work reported on here was designed to address directly these two potential objections. Specifically, the final study was based on the idea that the numerical probability information in standard classical paradigms can be replaced with equivalent information derived from participants' ability to predict events on the basis of experience. An example is Fox and Tversky's (1998) study in which they trained participants to predict fictional inflation and interest rate movements. The participants were later asked to choose between options for which the probability information had been replaced by inflation and interest rate information. Wu et al. (2009) used a similar methodology but instead replaced numerical probabilities with perceptuo-motor targets (that participants had been trained to hit).

As noted in Footnote 11, Wu et al. (2009) conducted two experiments. In one they tested whether the degree to which the independence axiom is violated differs across perceptuo-motor and classical decisions. They found that the axiom tended to be violated to similar extents for both types of decisions. In a second experiment they found that best fitting parameters of cumulative prospect theory (see Chapter 6 for a description) suggested that participants underweight low probabilities for perceptuo-motor decisions and overweight low probabilities classical decisions.

Interestingly, Fox and Tversky's (1998) results dissociate from those of Wu et al. (2009). Fox and Tversky's results suggest that choices based on learnt probabilities are equivalent to ones based on numerical probabilities, whereas Wu et al.'s (2009) do not. One possible explanation for this difference is that Fox and Tversky took into account peoples' subjective beliefs whilst Wu et al. assumed that the subjective probabilities matched the objective ones. This is a hypothesis that will receive indirect support in the next chapter.

The next chapter might also be viewed as relevant for another decision-making dissociation in the literature: the description-experience gap (Hertwig & Erev, 2009). The description-experience gap describes the apparent tendency for human choices to depend on whether options are experienced or described. A large cognitive literature (see Hertwig & Erev, 2009 and Rakow & Newell, 2010 for reviews) has grown from the finding that when outcomes of actions are experienced decisions show qualitatively different patterns compared to when outcomes are verbally described (given in numerical format). One potential shortcoming of such comparisons is that people have the opportunity to learn expected values directly in the decision from experience literature (unlike the studies of Fox & Tversky and Wu et al. cited above). This means that the argument presented in *General Introduction*, about how such choice situations

are radically different from the ones in the classical literature, applies with full force. Nevertheless, it is noteworthy that Wu et al. (2009) obtained the same probability weighting functions for their pointing task that have been reported repeatedly within the decision from experience literature (see e.g., Ungemach et al., 2009) – although Wu et al. themselves did not draw out this link.

In the next study, we, like Fox and Tversky (1998) and Wu et al. (2009), let our participants learn about their own task performance. Our participants completed both perceptuo-motor (pointing) and cognitive (mental arithmetic) learning tasks. In a separate decision session they made hypothetical choices between many pair-wise options. They did not receive any feedback on their choices and could therefore not learn a particular response strategy through feedback. In other words, the decision session was typical of those in the classical literature, with the exception that each participant made many such choices and one of the chosen options ways played out “for real” at the end of the experiment (i.e., choices were consequential). The next chapter describes the, to the best of our knowledge, first study to systematically apply both the type of performance standard applied in low-level decision studies and the performance standards applied in high-level decision studies, across three precisely matched low- and high-level decisions tasks.

6. The Gaps That Weren't: The Perception-Cognition & The Description-Experience Gap

As noted throughout this thesis, decades of research suggest that people do not maximize expected utility (e.g., Allais, 1952/1979; Ellsberg, 1961; Kahneman & Tversky, 1979; Tversky & Kahneman, 1981; 1992). Classical studies, which demonstrate sub-optimal decision making, typically ask participants to make hypothetical choices between pairs of options. You might be asked if you prefer option A1: “£4000 with a probability of .8”, or option B1: “£3000 with certainty”. Your next decision might be between option A2: “£4000 with a probability of .2” and option B2: “£3000 with a probability of .25”. The latter pair is derived by dividing each probability in the first pair by 4. If you are like most people, you prefer B1 and A2, and your preferences are inconsistent with expected utility theory (specifically, they violate the independence axiom, see Allais, 1952/1979; Kahneman & Tversky, 1979). By asking people to make such pair-wise choices, researchers have uncovered many departures from optimal decision making. More recently, however, two other decision making paradigms have produced results that seem to diverge dramatically from the classical ones.

The first of these paradigms will be highly familiar by now. In sharp contrast to the classical studies, recent studies of decision making in perceptuo-motor (Trommershäuser, Maloney & Landy, 2003a,2003b) and perceptual (Whitely & Sahani, 2008; Navalpakkam, Koch & Perona, 2009, Navalpakkam, Koch, Rangel & Perona, 2010) domains typically report optimal or near-optimal decisions. Like the classical studies, these studies provide participants with numerical *value* information. However, the numerical *probability* information is replaced by analogous information derived from lower level systems (see below for an example). Thus, it seems that low-level decisions based on *internal estimates of probability* are near-optimal, whereas high-level decisions based on *numerical probabilities* are sub-optimal: the *perception-cognition gap* (Trommershäuser, Landy & Maloney, 2006; Trommershäuser, Maloney & Landy, 2008).

Secondly, within more cognitive contexts, “decisions from experience” dissociate from the classical so-called “decisions from description”: the *description-experience gap* (Hertwig & Erev, 2009; Rakow & Newell, 2010). This gap is, unlike the perception-cognition gap, not about levels of performance. Instead this gap refers to people behaving as if they *overweight* low probabilities in the classical experiments, and as if they *underweight* low probabilities when relying on experience. In the latter

experiments, participants experience outcomes in the form of monetary rewards. Typically participants get to sample two computer buttons, each generating monetary outcomes with some probability. One button might return £32 with a probability of .1, and the other £3 with certainty (e.g., Hertwig, Barron, Weber & Erev, 2004). After sampling both buttons, participants choose which button to play for real money.¹⁹ Had participants been presented with the above options in numerical format (as in the classical studies) ~50% would have chosen the certain option. In contrast, when relying on experienced outcomes ~80% choose the certain option (Hertwig et al., 2004).

The underlying causes of both of these “gaps” have attracted considerable attention (Maloney et al., 2007; Hertwig & Erev, 2009; Rakow & Newell, 2010), but are far from resolved. Moreover, each “gap” has been pursued largely in isolation, but the relationship between the two literatures is of importance. Both literatures involve choice tasks that are based on implicit, or learnt, probability information (unlike the classical paradigm within decision-making research, but see Fox & Tversky, 1998) and both typically involve extensive feedback (again, unlike the classical paradigm).

Although, low-level decision making paradigms typically find optimal performance (e.g., Trommershäuser et al., 2008), and the decisions-from-experience paradigm finds deviations from the optimal strategy (Hertwig & Erev, 2009), recent results suggest that similar deviations from optimality may obtain in both paradigms (i.e., underweighting of low probabilities, compare Wu, Delgado & Maloney, 2009 and Hertwig & Erev, 2009). It is therefore tempting to conclude that a converging picture is emerging, namely, one in which decisions based on numerical probabilities dissociate from decisions based on internal estimates of probabilities (but see Fox & Tversky, 1998), and, furthermore, where tasks involving internal estimates behave in potentially similar ways, regardless of where in the cognitive system that internal estimate resides. However, as we argue next, confounds exist which suggest that such an inference may be premature.

Firstly, in contrast to classical studies (e.g., Kahneman & Tversky, 1979), participants are typically given the opportunity to *learn* how to choose. Some researchers interested in low-level decisions are aware of this, and present evidence against the use of learning strategies in their work (e.g., Trommershäuser et al. 2003a). In decision from experience studies, however, this is rarely considered a problem. In these studies, participants could potentially learn which of the two buttons to prefer by

¹⁹ Variants of this sampling paradigm in which each “sample” is played out for real with varying levels of feedback are also used, see Hertwig & Erev (2009). Note also that Hertwig et al. 2004 paid participants in € and not £.

tracking how much money each button returns on average (Hau, Pleskac, Kiefer & Hertwig, 2008). If such a tallying approach were taken, a direct trade-off between probabilities and values, as in the classical studies, would not be needed. Instead, the decision problem reduces to a choice between two expected values. The possibility that classical decision studies and the two more recent literatures investigate phenomena that are sub-served by functionally dissociable systems - a more simple learning system (available to humans, but also e.g., mice Thorndike, 1911) and a more complex system relying on representations of probabilities and values (as in classical studies) - suggests both that contrasts should be made with care and that “dissociations” may potentially be unsurprising.

Secondly, for the perception-cognition gap in particular, there is a real possibility that the gap is caused by methodological differences in study design (see e.g., Maloney et al., 2007) and/or differences in how performance is assessed. The latter possibility is rarely considered. Lower-level decision studies typically evaluate how closely people’s earnings match those of a hypothetical optimal agent (e.g., Trommershäuser et al., 2003). This contrasts with both classical and experience-based studies, which do not evaluate actual performance, but look for patterned deviations from optimality (e.g., Kahneman & Tversky, 1979; Hertwig & Erev, 2009). It is conceivable that the two ways of analysing performance produce different answers: one which characterizes human decision making as near-optimal, and one which characterizes it as sub-optimal. This result might obtain if, for example, patterned deviations, such as those found in classical studies, are not very costly.

The fundamental question we seek to explore is whether there really is robust evidence for “gaps”. To probe the issue, we compared three types of decisions under precisely matched conditions: classical decisions, decisions for which the numerical probability information was replaced with participants’ estimates of their performance on a low-level perceptual task (specifically, a standard pointing task) and decisions where the probability information involved estimates of participants’ own performance on a high-level task (specifically, a novel arithmetic task). In keeping with the classical studies, our design did not allow for an expected value learning strategy. To our knowledge, ours is also the first study applying both low- and high-level standards of performance to precisely matched tasks. This allowed us to test directly the idea that the perception-cognition gap might simply be due to the application of different standards.

To preview the results: As in the lower-level decision making studies our participants made highly efficient decisions. Crucially, performance was equally good

across low-level, high-level and classical tasks. When we fit cumulative prospect theory (Tversky & Kahneman, 1992), a model that has been used to account both for decisions from experience and decisions from description, we nevertheless detected patterned deviations from optimality. When we fit objective probabilities, the deviations were as predicted by the description-experience gap. However, participants' subjective beliefs did not match objective probabilities. When participants' biases were taken into account, decisions from experience showed the same systematic deviations as classical decisions typically do.

6.1 Methods

6.1.1 Participants and apparatus

Eighteen members of the School's participant panel were paid £6/hr (plus a possible bonus of £0-£6) to participate. Each participant took part in two learning sessions each lasting 45 minutes and one decision session lasting 1 hour. Ten participants came back approximately two weeks later for a repeat decision session.

Informed consent was obtained and the study was approved by the School's Ethics Committee. A Wacom tablet was used for the pointing task. Experiments were written in Matlab using Psychtoolbox (Brainard, 1997; Pelli, 1997; Kleiner, Brainard & Pelli, 2007).

6.1.2 Stimuli, design & procedure

The experiment involved two parts and made use of the idea that standard decision tasks can be altered to instead require participants to use their own internal estimates of probabilities (see Fox & Tversky, 1998; Wu et al., 2009) and was broadly based on Wu et al.'s Experiment 2. First, participants practised two tasks (mental arithmetic and pointing) in separate counterbalanced sessions. These sessions allowed participants to learn about their own task performance. The second part was a decision making session. It involved three tasks: a classical decision task with numerical probability information and two tasks in which the numerical probability information was replaced with *equivalent* low-level (perceptuo-motor) and high-level (mental arithmetic) information. In other words, participants saw exactly the same decision problems across all three tasks.

Participants' goal in the learning sessions was to earn as many points as possible. To get points, participants had to hit mental arithmetic and pointing targets under time

pressure (details below). Target hits were awarded 100 points, misses were not penalized and late responses cost -700 points. Participants were instructed to earn as many points as possible, but were told that learning about their own performance was at least as important as earning points. It was emphasized that improved knowledge would enable them to make better decisions in the third session, and that those decisions would have real financial implications.

The arithmetic task involved summing up four numbers (central numbers, Fig. 6.1A). Participants typed their response (to the nearest integer) on a virtual keypad (not shown here). Arithmetic targets were defined relative to the sum of the four presented numbers. A target of ± 6 (Fig 6.1A, B), for example, meant that the difference between the judged sum and the actual sum had to be smaller than ± 6 to count as a “hit”. After having responded participants received feedback on how far their judgement was from the actual sum (Fig. 6B).

The pointing session involved pointing towards targets in the frontal plane and was based on Wu et al.’s (2009) pointing task. To score a hit, participants had to hit anywhere on the shaded bar in Fig. 6.1 D (illustrating a pointing target). Pointing targets were displayed in random locations and made sufficiently tall so that only variability along the x-axis mattered. Feedback consisted of a high-contrast disc showing where the screen was hit (Fig. 1E). To encourage participants they also received explicit “hit” and “miss” feedback after each trial (Fig 6.1B & E) for both tasks.

The target centre in the pointing task was the midline of targets. The target centre in the arithmetic task was the sum. Because time was limited, participants were not able to perform either task with full accuracy. Subtracting the target centre from each response results in error distributions (Fig. 6.1 C, F) describing the accuracy and precision of individual participants’ responses. These can be used to predict participants’ chances of hitting targets, both pointing and numerical, of varying “widths”.

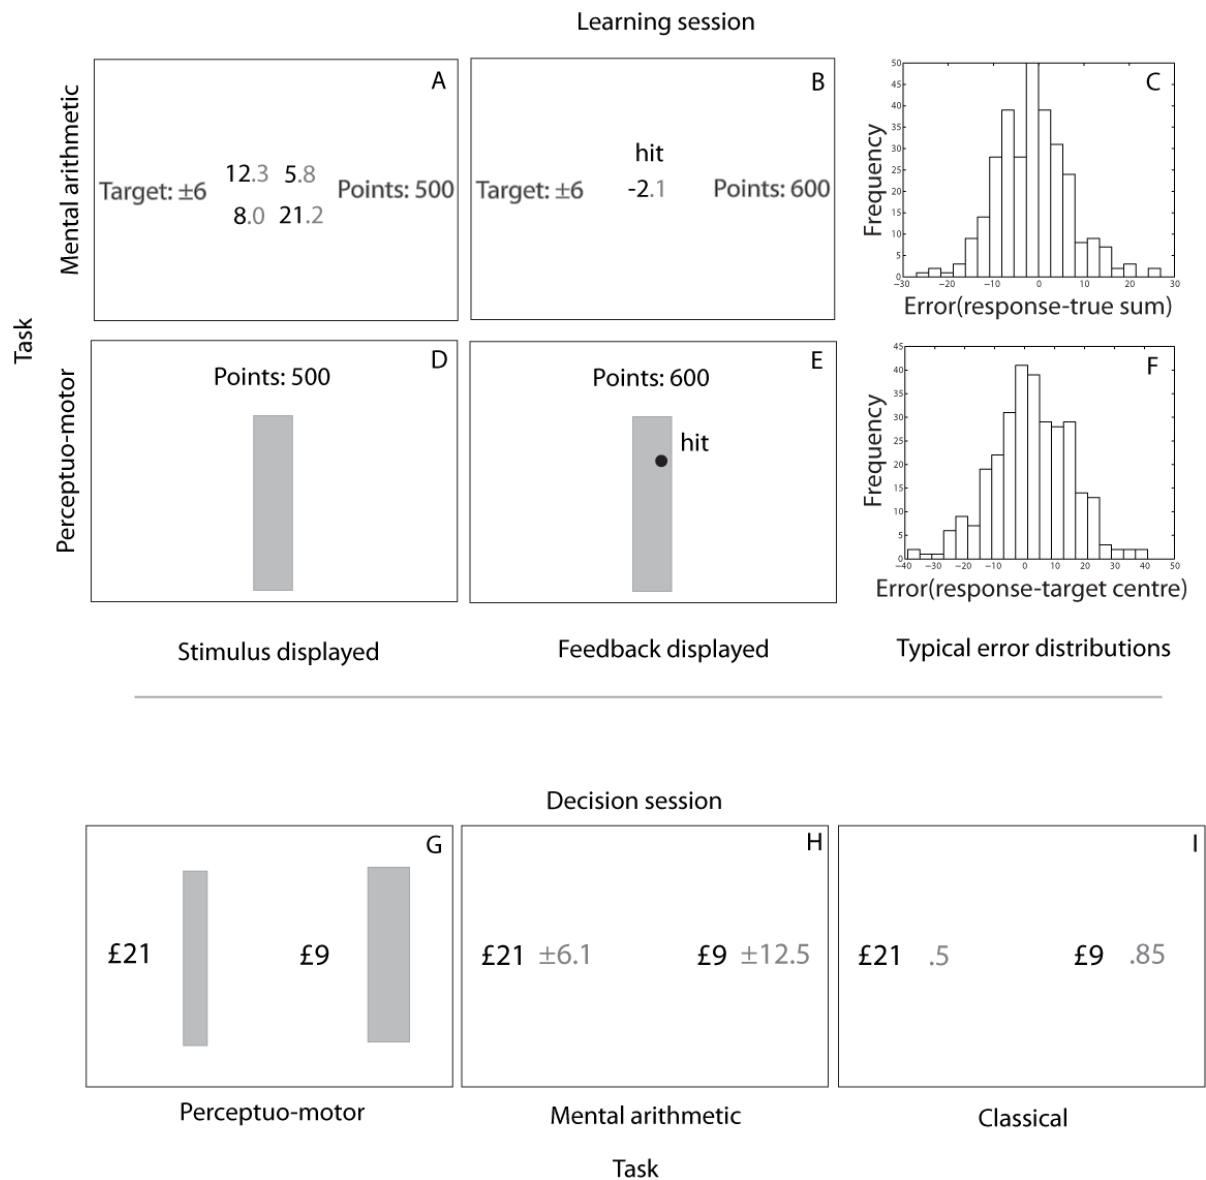


Fig. 6.1. Experimental paradigm. Panels A-F illustrate learning sessions. Panel A: Mental arithmetic stimuli. The task involved summing the four central numbers. The judged sum had to be within the limits of the target (± 6 here) to be scored a hit. The total number of points earned was displayed to the right. Panel B: Mental arithmetic feedback. Feedback included the error (difference between the judged and true sum, -2.1 here) and explicit hit/miss information. Panel C: An arithmetic error distribution (Participant 10). Panel D: Pointing stimuli. The bar represents the target and “Points 500” illustrates the cumulative score information. Panel E: Pointing feedback included explicit hit/miss information as well as a high-contrast indication of where the screen was hit (black disc here). Panel F: A pointing error distribution (Participant 10). Panel G-I illustrate the decision session. Panel G: two choice options with probabilities replaced by pointing targets. Panel H: two choice options with probabilities replaced by arithmetic targets. Panel I: two classical choice options with numerical probabilities.

The decision session involved choices between pairs of options (Fig 6.1G-I). Each option contained probability information in one of three formats: low-level (pointing, Fig. 6.1G), high-level (mental-arithmetic, Fig. 6.1H) and classical (numerical, Fig. 6.1I) as well as value information. There were 120 option-pairs for each format presented in a randomized order with no time limit. For each pair, participants indicated whether they preferred the left or the right option. Importantly, pairs of options were *matched* across the tasks. That is, each classical option had an exact equivalent in the cognitive and the perceptuo-motor domain (target widths were adjusted such that target hit probabilities matched classical probabilities for each participant).

Target hit probabilities were matched to classical probabilities as follows: Target widths corresponding to objective hit probabilities were estimated from response distributions (as illustrated in Fig. 6.1 C & F, 300 data points per participant and task) by fitting Gaussians and using numerical methods to find the target width corresponding to each reference probability (integrating the fit Gaussian over target widths). For all but one participant, who showed a quite peaked distribution for the mental arithmetic task, the normality assumption was well met.

Unlike most previous studies (e.g., Kahneman & Tversky, 1979; Ungemach et al., 2009; Wu et al. 2009; Camilleri & Newell, 2011, but see Erev, Roth, Slonim & Barron, 2002), we used randomly selected choice options. Probabilities were drawn from the range .05 to .95 (in steps of .05). Values were drawn from £1 and £3 to £54 (in steps of £3). This resulted in options with a wide range of differences in expected values.

Of the 120 option-pairs, 5 were pairs in which one option dominated the other (i.e., both probability and value was higher for one option). Most participants chose the dominating option most of the time (mean number out of 5 = 4.85, min=4, max=5), indicating that participants paid attention and understood the decision task.

Participants did not receive feedback but knew that one of their chosen options would be randomly selected at the end (with values decreased by a tenth) and played for real money. They also knew that the probabilities used to generate the real outcome would be based on their own performance in the learning sessions.

6.2 Results & discussion

6.2.1 How good were people's choices?

We first sought to determine how good participants' choices were using metrics commonly employed in low-level decision studies. The perception-cognition gap

implies that people will make much better choices when they are faced with low-level decisions compared to when faced with classical decisions, or compared to when faced with decisions based on high-level information.

To achieve results comparable to low-level studies, and because expected value maximization might be considered normative in typical decision making studies (Rabin, 2000, Rabin & Thaler, 2001), we evaluated how close participants came to maximizing expected value. An expected value maximizer chooses the option that will return the most reward in the long run. As expected value maximization does not allow weighting of objective values it is a stricter normative standard than expected utility maximization.




Fig. 6.2. Performance metrics as a function of task. Panel A. The proportion of choices maximizing expected value for each participant (gray discs) and the group average (black squares). Error bars are parametric 95% confidence intervals. Panel B. The proportion of choices maximizing expected value for each task as a function of the difference in expected value of the choices options (discriminability) pooled across participants. Panel C. Efficiency for each participant (gray discs) and the group average (black squares). Error bars are parametric 95% confidence intervals. Gray discs representing individuals have been jittered laterally.

Fig. 6.2A shows the proportion of optimal choices both for individual participants (grey discs) and the average proportion of optimal choices for each task (black squares). As can be seen, the average proportion of choices that maximized expected value was approximately .75 - regardless of task. Note also that there were large individual differences.

An average optimal choice rate of .75 is moderately impressive. However, expected value maximization implies that the decision maker is able to perfectly discriminate between choice options (as does expected utility theory). It is implausible that humans can achieve perfect discrimination. Indeed, previous studies have found that choice consistency increases when choice options become more discriminable (i.e., when the difference in utility between choice options increases, see e.g., Mosteller & Nogee, 1951).

As Fig. 6.2B shows, the proportion of optimal choices amongst our participants increased as choice options became more discriminable (i.e., the expected value difference increased). Note, increasing differences between choice options does not only affect people's ability to discriminate, it is also related to the potential loss of making an incorrect choice. The easier it is to discriminate between options the more costly mistakes become, and conversely the harder the discrimination the less costly the mistakes. If participants choose the wrong option mainly when options are hard to discriminate Fig. 6.2A may give an overly pessimistic picture of participants' choice performance.

Fig. 6.2C shows our participants expected earnings relative to a hypothetical participant who always chooses optimally (i.e., relative to someone who always chooses the option with the highest expected value). Efficiency, or actual gains over gains achieved by an optimal participant, is the standard performance metric in lower-level decision studies (e.g., Trommershäuser et al., 2003a). As can be seen, the average participant is expected to earn ~92% of the optimal earnings (with some expected to earn nearer to 98%). Thus, whatever choice strategies our participants used – they were nearly as efficient as the optimal one.

Fig. 6.2C also suggests that there were next to no differences in performance across the three tasks. For statistical comparisons when the null hypothesis is of interest Bayesian tests are appropriate (Gallistel 2009, Rouder et al., 2009). A comparison using JZS-Bayes factors (Roeder et al., 2009) supports the hypotheses of equal efficiency levels across the three tasks (Cognitive Vs Pointing: Bayes Factor = 5.59, $t(17) = -.249$, $p = .81$; Cognitive Vs Classical: Bayes Factor = 5.3, $t(17) = .335$, $p = .74$; Pointing vs. Classical: Bayes Factor = 3.56, $t(17) = .979$, $p = .34$).

Thus, we found evidence *against* the perception-cognition gap. Performance was equally good across low- and high-level decisions. Moreover, because performance was just as good in the classical task, the results also imply an upper bound on the description-experience gap. That is, although equal performance does not imply equal

process, equal performance implies that any differences in process are inconsequential behaviourally.

6.2.2 Do people deviate systematically from optimality?

We next sought to determine whether, despite the equally good performance, differences in choice strategies could be detected. Specifically, we sought to determine whether there were differences in how probabilities were treated across the three tasks - as predicted by the description-experience gap. Any decision model that allows for under and overweighting of probabilities could in principle be used to test for differences in probability weighting.

We fit cumulative prospect theory (Tversky & Kahneman, 1992), a model commonly used to account for deviations from optimal decision making. Two key aspects of the theory are its value and probability weighting functions. These functions map objective quantities onto subjective quantities, and so allow subjective values and probabilities to deviate systematically from objective ones. These aspects allow it to account for many of the empirically observed deviations from optimal choice.

Briefly, we used the parameterization recommended in Stott (2006): a Prelec (1998) one-parameter probability weighting function, a power value weighting function and a one-parameter logistic choice function. The choice function captures the fact that people are less-than-perfectly sensitive to differences between options (see Fig 6.2 B). The model was fit separately to each participant and task minimizing the log-likelihood. For additional details and a model-exploration using seven other parameterizations see *SI Additional methods: Model fitting*.

Because the description-experience gap manifests itself in differences in probability weighting, and because value weighting is generally not considered sub-optimal (but see Rabin, 2000), we focus on the best-fit probability weights in the following.

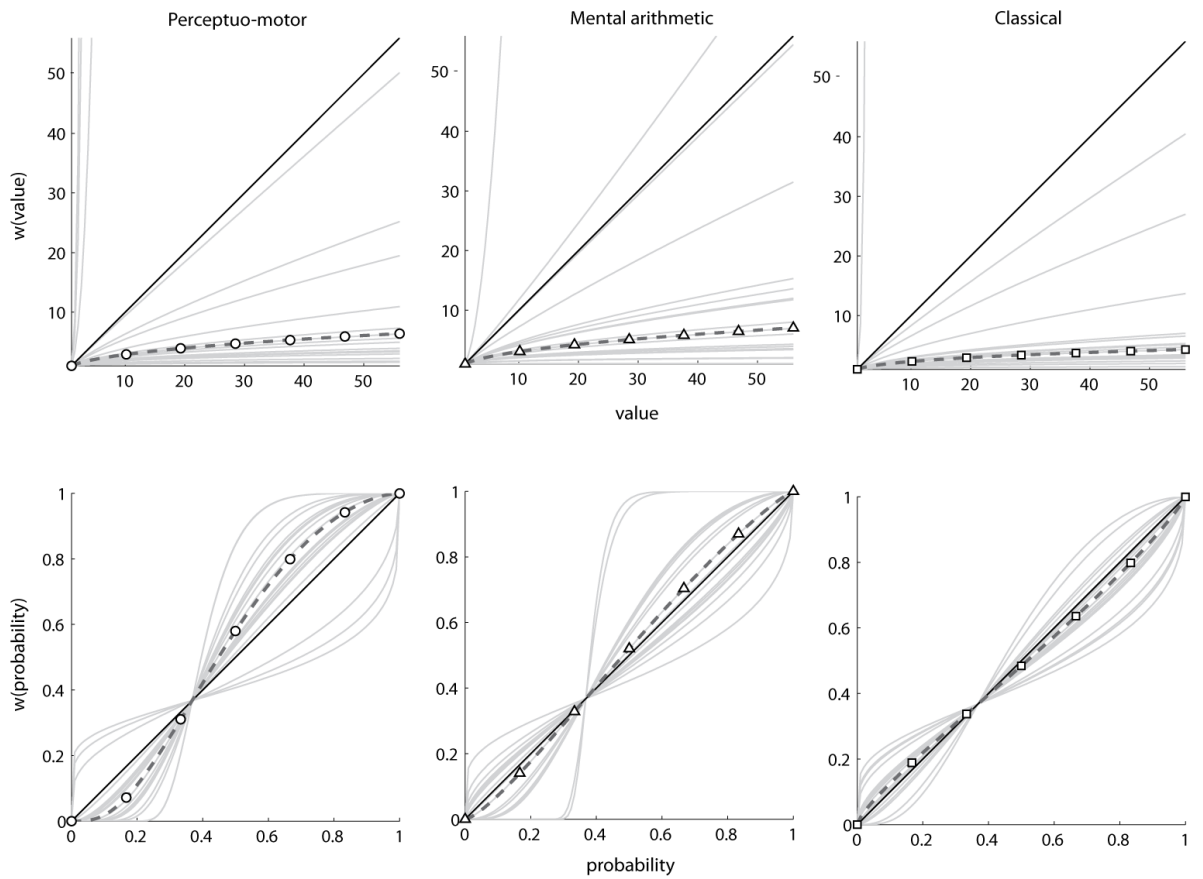


Fig. 6.3. Best fit value and probability weights. Row 1, Column 1-3: Participants' best fit value weights (gray lines) and group averages (black dashed lines)²⁰. Row 2, Column 1-3: Participants' best fit probability weights (gray lines) and group averages (black dashed lines). Circles = perceptuo-motor, triangles = mental arithmetic and squares = classical.

Fig. 6.3 shows the best fit value (row 1) and probability (row 2) weighting functions for each task and participant. Two trends are noteworthy. Firstly, the average probability weighting function (dashed lines, row 2), suggests underweighting of low probabilities for the pointing and arithmetic task, but suggests overweighting for decisions from description (i.e., “classical”) – thus seemingly replicating the description-experience gap (Hertwig & Erev, 2009)²¹ and replicating previous results comparing perceptuo-motor decisions to classical ones (Wu et al., 2009; 2011).

²⁰ Group averages are exponents of the 50% trimmed mean on the logarithm of individual weights. This method accounts for the fact that the weights have qualitative different meanings in the range 0-1 and in the range 1 to infinity.

²¹ Note that although underweighting of rare events is considered a characteristic of decisions from experience (Hertwig & Erev, 2009; Rakow & Newell, 2010) the picture emerging from studies that fit (cumulative) prospect theory in order to evaluate underweighting, instead of inferring underweighting from peoples' choice patterns (e.g., Hertwig et al., 2004), is somewhat mixed. Whilst some studies find probability weighting parameters suggestive of underweighting (e.g., Ungemach et al., 2009; Camilleri & Newell, 2011) others find weights suggesting only very marginal (i.e., near-linear) underweighting (e.g.,

The second noteworthy trend is that, consistent with previous studies (Wu & Gonzales, 1998; Wu et al., 2009; 2011), there were large individual differences. Some participants appear near-optimal with near-linear probability weights, whereas others show severe under- or overweighting. Specifically, for neither decision from experience, nor for decisions from description, is there a consistent pattern of only under- or overweighting.

If these individual differences are consistent and not, for example, due to noise, using a given participant's best-fit parameters for one task (e.g., classical), to predict their choices in another task (e.g., pointing) should yield better predictions, compared to predicting the same participant's choices using the average best-fit parameters for one and the same task (e.g., classical).

To test this, we predicted each participant's choices using the average best-fit parameters for the same task (exponent of the 50% trimmed mean [excluding the predicted participant] log parameter value). This resulted in an 18x3 matrix of log-likelihoods. We summed across tasks to create an aggregate measure of within-task-between-subject predictability for each of the 18 participants.

We also predicted each participant's choices for each task using their best-fit parameters for each other task. This resulted in 2 predictions for each task and each participant (e.g., predicting classical from arithmetic and pointing). We averaged over the two predictions for each task and summed over tasks producing, as for the within-task case, a vector with 18 data points for between-task-within-subject predictability.

We compared the vectors using a paired t-test having removed three outliers (see SI Fig. S6.1). The test shows that if we want to predict your choices for a particular task, it is better to measure your responses for an *alternative* task, compared to trying to predict your responses on the basis of other people's choices for the *same* task ($t(14) = 5.66, p < .0001$). This suggests that the differences in average probability weights across tasks (dashed lines, Fig. 6.3, row 2) are of limited importance.

6.2.3 Are peoples subjective beliefs calibrated?

Hau et al., 2008), whilst others find overweighting (e.g., Fox & Hadar, 2006; Abdellaoui, L'Haridon, & Paraschiv, 2011). Without systematic study it is difficult to trace the origin of these differences. The studies often differ substantially both in design and analysis methods. However, note that we use a large number of different choice options, fit individual choice data in a setting where expected value learning strategies have been ruled out and use a maximum-likelihood approach which essentially avoids the problem of flat maxima (c.f., e.g., Ungemach et al., 2009). Importantly, we replicate studies employing a similarly powerful design and analysis methods (see e.g., Wu et al., 2009).

In the learning sessions our participants were not only given the chance to learn about their task performance, but were also explicitly asked about their beliefs about their performance. Every 50th trial (300 in total), we asked participants to adjust target widths so that they thought they would hit them 95%, 75%, 50%, 25% and 5% of the time. If participants' subjective probabilities match objective ones they would, for example, upon being asked for a 50% target width, report a target width that would allow them to hit the target ~50% of the time. If they do this, they are calibrated. If they set a target width that is too wide, or too narrow, they are not calibrated and their beliefs do not (fully) match reality.

The previous model fitting assumed that people's internal beliefs about uncertainties match objective ones. Some studies have assessed whether their participants were calibrated as *a group* and often (e.g., Hau, Pleskac, Kiefer & Hertwig, 2008; Ungemach, Chater & Stewart, 2009; Gottlieb, Weiss & Chapman, 2007; Wu et al., 2009), but not always (e.g., Fox & Tversky, 1998; Wu, Delgado & Maloney, 2011), they are.

A group, however, might be perfectly calibrated on average, yet have every member showing substantial biases (Fig. S6.2). In other words, to assess calibration of individuals, one needs to look at individual calibration. Moreover, the task used here required participants to use general task knowledge to derive probability information. This is very different from tracking frequencies of specific events (e.g., how many times out of 40 one obtained £3 when pressing a particular button, Ungemach et al., 2009). Either, or both of these factors, may be responsible for the relatively good calibration evidenced in some previous studies.

Fig. 6.3A-D shows the relationship between subjective width ratings (gray symbols) and objective widths (white symbols) for two representative participants. Participant 16 (Row 1), who represents the most common pattern, generally overestimated target widths required to match the reference hit probabilities. Participant 12's (Row 2) judgments show a less common pattern best described as regressive (see Fig. S6.3 for plots of all participants' judgments).

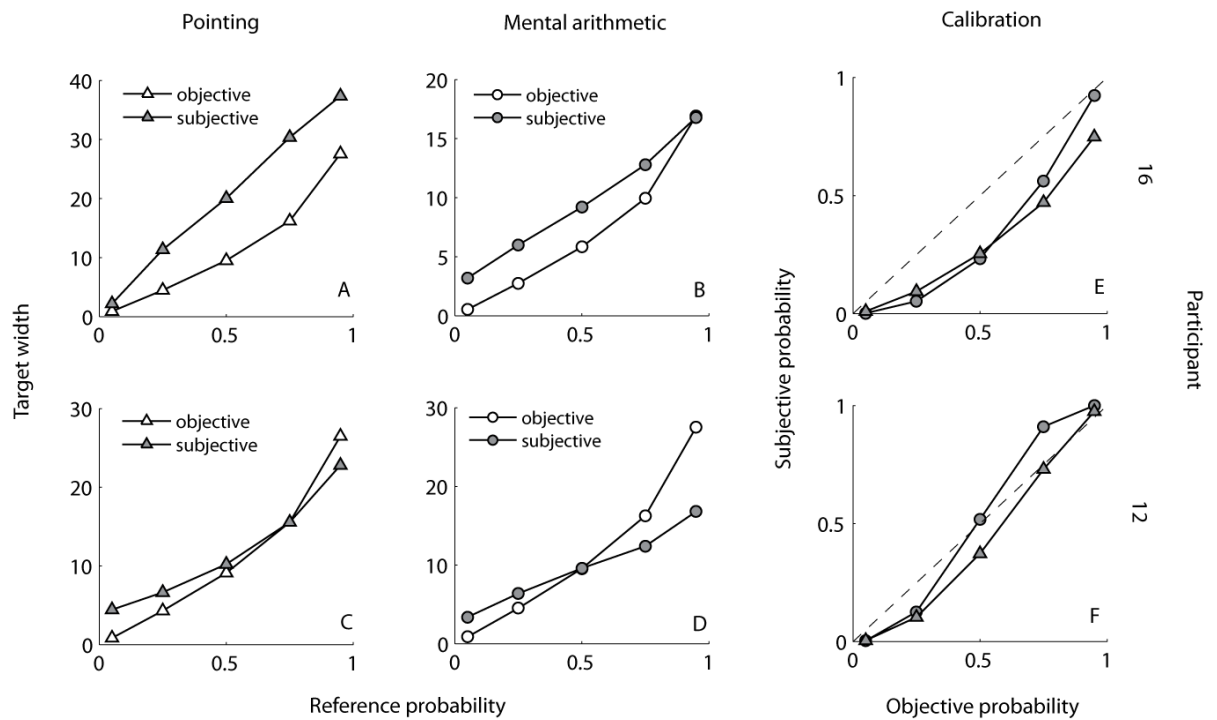


Fig 6.4. Calibration of two representative participants. Row 1 = Participant 16, Row 2 = Participant 12. Panel A-D. Target widths matching reference probabilities (objective), and target widths judged to match reference probabilities (subjective = average of the last 5 of 6 width judgments). Panel E&F. Subjective hit probabilities as a function of objective hit probabilities.

Fig. 6.4E and F shows these same biases expressed as participants' belief in their ability to hit targets as a function of their actual ability. Participant 16 generally underestimated their ability, whilst Participant 12 underestimated their ability for hard targets and overestimated their ability for easy targets. Note also the intra-subject consistency in biases across tasks. When averaged across participants, the group bias (Fig. S6.4) is similar to Participant 16's and to some previous group-average results (Wu et al., 2011).

6.2.4 Do people deviate from optimality when biases are taken into account?

The description-experience gap is typically evaluated under the assumption that objective probabilities correspond to subjective probabilities (see e.g., Camilleri & Newell, 2011, and perceptuo-motor studies have assumed the same (Wu et al. 2009; 2011). If subjective probabilities are not calibrated, biases and not underlying changes in preferences may underlie the gaps. This is particularly relevant as the average participant showed a general underestimation of probabilities (Fig. S6.3, see also Wu et

al., 2011). Perhaps, underweighting in decision from experience studies can be accounted for by this bias?

To test this idea, we repeated the model fitting having replaced the objective probabilities for each choice option with probabilities estimated from participants' subjective ratings. We used Weibull functions to extrapolate from subjects' width ratings subjective probabilities other than those corresponding to the reference probabilities. As can be seen in Fig. S6.3, Weibull functions generally capture the width ratings well.

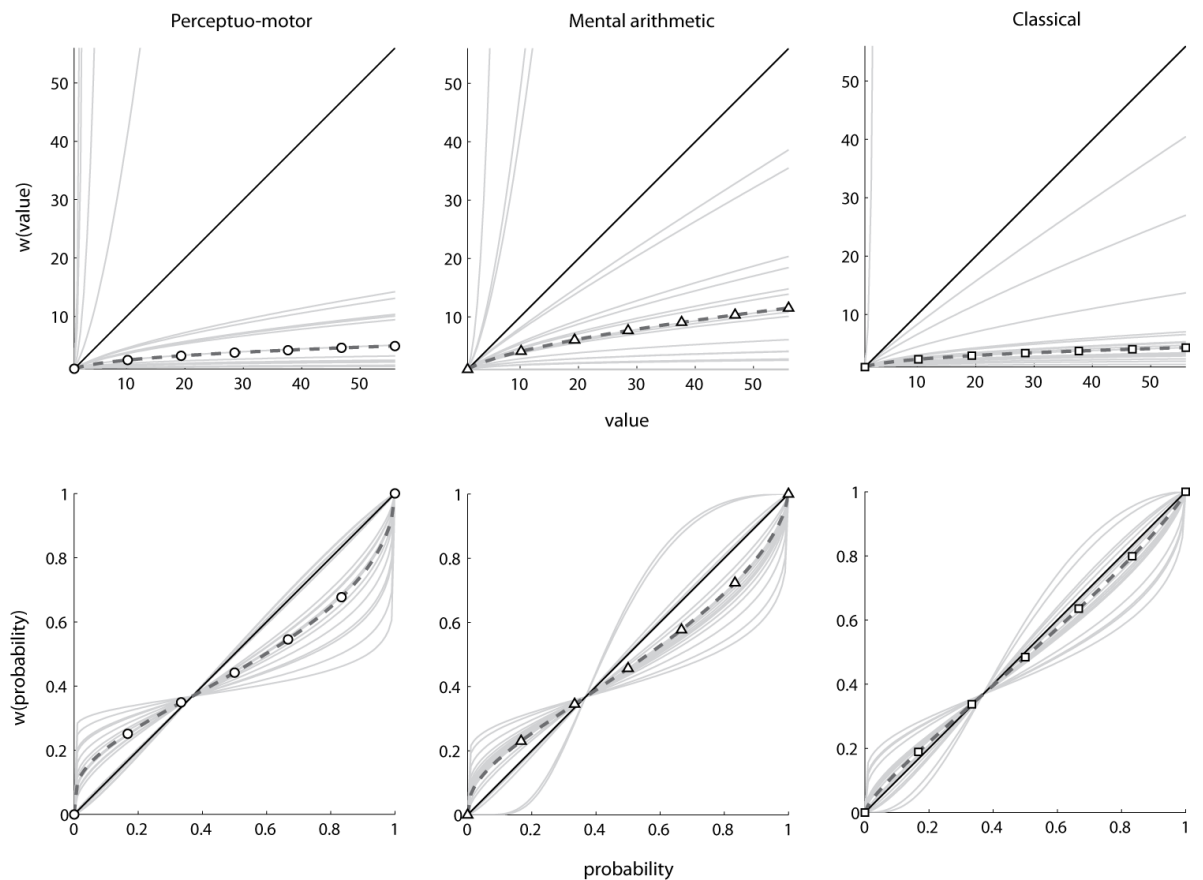


Fig. 6.5. Best fit probability weight parameters (based on subjective probabilities). Row 1: Participants' best fit probability weights (gray lines) and group averages (black dashed lines). Circles = perceptuo-motor, triangles = mental arithmetic and squares = classical. Group averages are exponents of the 50% trimmed mean on the logarithm of individual probability weights.

When subjective rather than objective probabilities are fit the results change dramatically (Fig. 6.5). Now, most participants appear to overweight small probabilities – just as they do when probabilities are given explicitly. Of course, the above fits (Fig. 6.5 & Fig. 6.3) are due to one particular parameterization of cumulative prospect theory

(the one recommended in Stott, 2006). Thus, it is possible that the above results obtain only with this particular parameterization. The space of possible parameterizations is large (Stott alone tested 256 different parameterizations). It is not feasible to search the entire model space, but we nevertheless sought to verify that the results above do not depend on choosing a particular model instantiation. We provide details on this test in “SI Additional model fits” and note here that a comparison across eight different parameterizations produces results concordant with Fig. 6.5 and Fig. 6.3.

Another potential caveat is the relationship between reported probability judgments and choices. It is possible that our participants were simply bad at reporting their beliefs. That is, the probabilities participants reported and the probabilities they based their choices on might dissociate. This does not appear to be the case. A simple model, assuming expected value maximization and less-than-perfect discrimination, generally better accounts for the data when subjective probabilities are fit, relative to when objective probabilities are fit (Fig. S6.5) – indicating that subjective probability estimates provide meaningful information over and above objective probabilities.

6.3 Summary

We compared decisions based on three different types of probability information, low-level (perceptuo-motor), high-level (mental arithmetic) and numerical, under identical conditions. We found evidence against the perception-cognition gap. Our participants equally made highly efficient decisions when relying upon perceptuo-motor knowledge, upon cognitive knowledge or on numerical probabilities. We nevertheless replicated the description-experience gap when fitting objective probabilities. However, subjective probabilities did not match objective probabilities. When people’s biases were taken into account, there was no gap. In fact, when biases were taken into account, both decisions from experience (whether cognitive or perceptuo-motor) and classical decisions showed the same overweighting of small probabilities.

What factors might explain the discrepancy between our results and previous results in support of decision making “gaps”? Arguably, the most parsimonious explanation for the perception-cognition gap lies in the very different methods by which performance has been assessed. Indeed, when we used the standard low-level methods, that of evaluating actual choice performance, we found that although participants were not precisely optimal, they were nevertheless highly efficient. In contrast, when we fit cumulative prospect theory, many participants had probability weights that were substantially different from 1, suggesting that they sub-optimally weight probabilities.

Thus, we were able to show both “good” performance (using standard low-level metrics) and “bad” performance (using standard high-level metrics) for *both* low- and high-level decisions.

The reason for the absence of a description-experience gap is perhaps less clear. We replicated the gap with objective probabilities, but not when fitting subjective probabilities. This might be taken to imply that the gap has arisen because it is erroneously assumed that subjective probabilities match objective ones. However, as noted in the introduction, standard decision-from-experience paradigms have a potential learning component that might also explain the gap. In that paradigm, participants also learn frequencies and do have to use task-relevant knowledge to estimate probabilities as our participants did (the former is knowledge specific to a particular button, the latter involves generalisation from task knowledge to any number of possible targets). Any of these factors might account for the different results. The potential role of learning in decision-from-experience tasks, however, means that the decision task used here is closer to the classical tasks. It might therefore be argued that the present study represents a better test of the description-experience gap.

Incidentally, there is much current debate about whether people have stable preferences (see e.g. a recent special issue: Schaik, Kusev & Juliusson, 2011). If they do not, fitting models assuming that they do (such as cumulative prospect theory or expected utility theory) is non-sensical. Two weeks after the decision session we asked our participants if they would come back for another session and 10 did. We correlated each participant’s 120 choices in Session 1 with their choices in Session 2 (the option-pairs were presented in a different random order). The choices were highly correlated across the two sessions (mean $r = .6$ for pointing, $r = .58$ for arithmetic, $r = .64$ for classical, 54 p -values $< .0001$, 1 p -value = $.018$). For each of the three tasks, three participants reported changing their choice strategy across the two sessions. When these participants are excluded the mean correlations increase substantially (mean $r = .69$ for pointing, $r = .68$ for arithmetic, $r = .71$ for classical). This shows that past choices predict future choices and suggests that people know something about their own choice consistency.

Despite equally good performance across tasks, despite the same apparent sub-optimal weighting of probabilities across tasks, and despite the relative stability of preferences across time, our results do *not* imply that decisions made using different modalities, or on the basis of different kinds of information, are *identical*. It is, for example, conceivable that decisions can be shown to differ across modalities when tasks

are tweaked specifically for this purpose. Moreover, our participants made perceptuo-motor decisions faster than they made either mental arithmetic or classical decisions (see Fig. S6.6). One might speculate that this is because probability information in the guise of perceptuo-motor targets is easier to discriminate. Either way, the fact that perceptuo-motor decisions were faster suggests that the processes underlying the decisions are not identical. Nevertheless, our results question whether such differences merit the distinction “gap”. In the most important way – that of actual earnings – there was no difference across the tasks. Importantly, because efficiencies were high, the ways in which people *do* deviate from optimal decision making do not seem particularly costly. Put another way, people may not be perfectly rational, but their irrationalities do not seem to lead them far astray from optimality.

7. General Discussion

7.1 Summary of empirical work

We began by reporting on a perceptuo-motor study (Chapter 2), in which we showed that optimality standards, commonly employed by those studying lower-level decision-making, are not absolute but relative. Performance standards are conditional upon the specifics of the tasks employed and the assumptions included in ones model. This meant that we could not, as first intended, use such ideal observer models as absolute standards. This, in turn, meant that we could not (without complication) use ideal observers to compare performance across perceptual and cognitive tasks. Instead, we had to find a way to design tasks such that perceptual and cognitive ability could be fairly compared.

We did this in the next series of experiments by evaluating how good people were at making decisions about how much time to spend on the task at hand (Chapter 4). Because the task at hand can be any task, we could compare peoples timing decisions, with no feedback and abstract reward structures, when the underlying task was either perceptual or of a more cognitive nature. In other words, we designed experiments for which the decision task and the reward information were *identical* across “modalities”, but for which the knowledge that decisions were based on was either perceptual or cognitive.

We found that our participants’ timing decisions were near-optimal, *whether* they were making decisions on the basis of *perceptual or cognitive information*. They were able to take both task difficulty and the reward structure into account. We even demonstrated that the ability to take task difficulty into account was dynamical, that is, could be made on a trial-by-trial basis. However, because only decisions about time were studied, the possibility remained that the good performance across “modalities” was due to timing decisions being special. Another aspect of this paradigm was that it provided no direct comparison of decisions based on numerical probabilities and probabilities derived through low-level experience (or for that matter high-level experience). In other words, not only was the type of decision restricted to the time domain, Chapter 4 also left open the possibility that decisions based on internal estimates of probabilities are optimal (whether cognitive or perceptual), whereas decisions based on numerical probabilities are not (but see e.g., Hertwig & Erev, 2009).

The final study we reported on (Chapter 6) was designed to address these potential issues. Here participants first learned about their own task performance for a perceptuo-motor and for a cognitive task. They then had to make decisions on the basis

of this task knowledge. Specifically, they had to make three types of decisions: decisions based on numerical probabilities, and decisions for which numerical probabilities and been replaced with the equivalent perceptuo- motor and cognitive information. Under such conditions, choice efficiencies were equally high across the three types of decisions. Thus, when decision tasks have been carefully matched, perceptual and cognitive performance appears equally good – suggesting that the perception-cognition gap is illusory.

As was pointed out throughout, there are many differences across low-level and classical decision-making paradigms. The former tends to, for example, include feedback and real payoffs. Any or all of the differences could potentially explain the perception-cognition gap (as observed elsewhere). Of course, given that we observed no gap, we might appear none the wiser as to the true cause of the “gap”. However, in Chapter 6, we also showed that participants appeared to *sub-optimally weight probabilities across both low- and high-level tasks*. This sub-optimality implies that the most likely explanation for the perception-cognition gap lies in the use of different performance standards across different literatures.

Thus, the perception-cognition gap probably arose because people were comparing *actual performance* in low-level tasks to whether or not *deviations from optimality* could be detected in high-level tasks. In other words, the apparent perception-cognition gap has arisen because people contrasted different types of studies asking different types of questions.

In summary, whether or not we view people’s decisions as good or bad seems to depend crucially on the performance criteria we apply. It appears that only when performance criteria and decision tasks are mismatched can one create the illusion of a gap. As long as one consistently applies the same performance criteria across matched tasks one finds very little evidence for a perception-cognition gap.

7.2 The optimality of human decision-making

What, if anything, do our results imply for questions about optimality and rationality in general? Researchers interested in questions about human rationality seem to slot neatly into two categories. Some believe that humans are optimal/rational and others that humans are sub-optimal/irrational. Both categories of researchers use empirical studies to support their arguments. Generally, the studies purport to show either surprisingly good performance, or surprisingly bad performance. We are to be

impressed or shocked by how good or bad decision makers humans are (for a particularly heated debate see Kahneman & Tversky, 1996; Gigerenzer, 1996).

Nevertheless, a third perhaps slightly less exciting story is conceivable. Perhaps humans are neither precisely rational nor precisely irrational. In Chapter 2, 4 and 6 we consistently found highly efficient choices. Yet, in the same chapters we also showed that deviations from optimal strategies were detectable. These results might be taken to suggest that human choice fairly closely *approximates* the optimal solution, yet is not fully optimal.

The efficiency metric is interesting precisely because it produces a gradient of rationality – not a binary optimal/sub-optimal categorization. Across the three chapters efficiencies have generally been in the .9 to 1 range. Clearly, one needs to compare efficiencies across studies with care. For example, the efficiency metric in Chapter 4 (time study) was such that an efficiency of 1 was the maximum achievable efficiency. In contrast, in Chapter 2 the optimal efficiency was distributed around 1, and greater efficiency than 1 was achievable. Nevertheless, the high efficiency across three quite different decision paradigms, in which the decision itself was quite different (Chapter 2– target/aim point choice; Chapter 4 – timing decisions; Chapter 6 – classical binary choices between pairs of options) might suggest that human choice in general is quite good.

There are two caveats to this story. One is that we have studied human behaviour in fairly constrained and simple laboratory tasks. Thus, we have not shown that human choice approximates optimal solutions more generally. One of the complaints against applying optimal standards (here expected utility theory) is that they do not scale very well (Gigerenzer, 2008). Thus, it is possible, for example, that highly efficient behaviour breaks down in more complex situations. On the other hand, our constrained and simple tasks employed arbitrary and relatively artificial reward structures (i.e., cost functions) making them in some ways more difficult than “everyday” tasks.

Clearly, for some tasks “optimal” solutions are implausible. For example, having a completely pair-wise consistent belief structure is computationally intractable (Nickerson, 2008). However, the mere fact that such tasks are implausible suggests that the rationality standards themselves are inadequate. If it would take one longer than one’s life time to do one consistency check across all one’s beliefs – then clearly this is not something one should do. From a decision theoretic perspective, the cost of making those checks will make them irrational. That is the cost function constrains what should be considered rational or not.

The other caveat is that whilst one might argue that ~90 – 95% absolute efficiency is rather good, clearly, if your stocks could return 100% but through sub-optimal management only returned 95% you might not be entirely happy (perhaps more so if 100% was break even and 95% was a 5% yearly loss).

Taken at face value however, the highly efficient behaviour of our participants raises questions about the processes underlying their behaviour. To what extent can people's behaviour be described as a result due to as-good-as-it-gets engineering solutions to difficult problems? Whether such models, or models based on sub-optimal solutions to hard problems (Gigerenzer, 2008), best account for human behaviour is arguably an open question.

7.3 Future directions

Both imaging (e.g., fMRI and MEG) and computational modelling seems to be on the rise in the study of the human mind. The relatively new “field” of neuroeconomics (Glimcher, Camerer, Fehr & Poldrack, 2009) is but one example of this. New ways of analysing imaging data and new imaging technologies together with the wider use of modelling raises the interesting possibility that advances will be made because behavioural data, models and imaging data can act as mutual constraints - each constraining the other in order to create better theories.

8. Supplemental materials

8.1 Chapter 2 – Supplementary materials

8.1.1 Movement & reaction time analyses

Here we break down the response times into its two separate components: reaction time and movement time and analyse them separately.

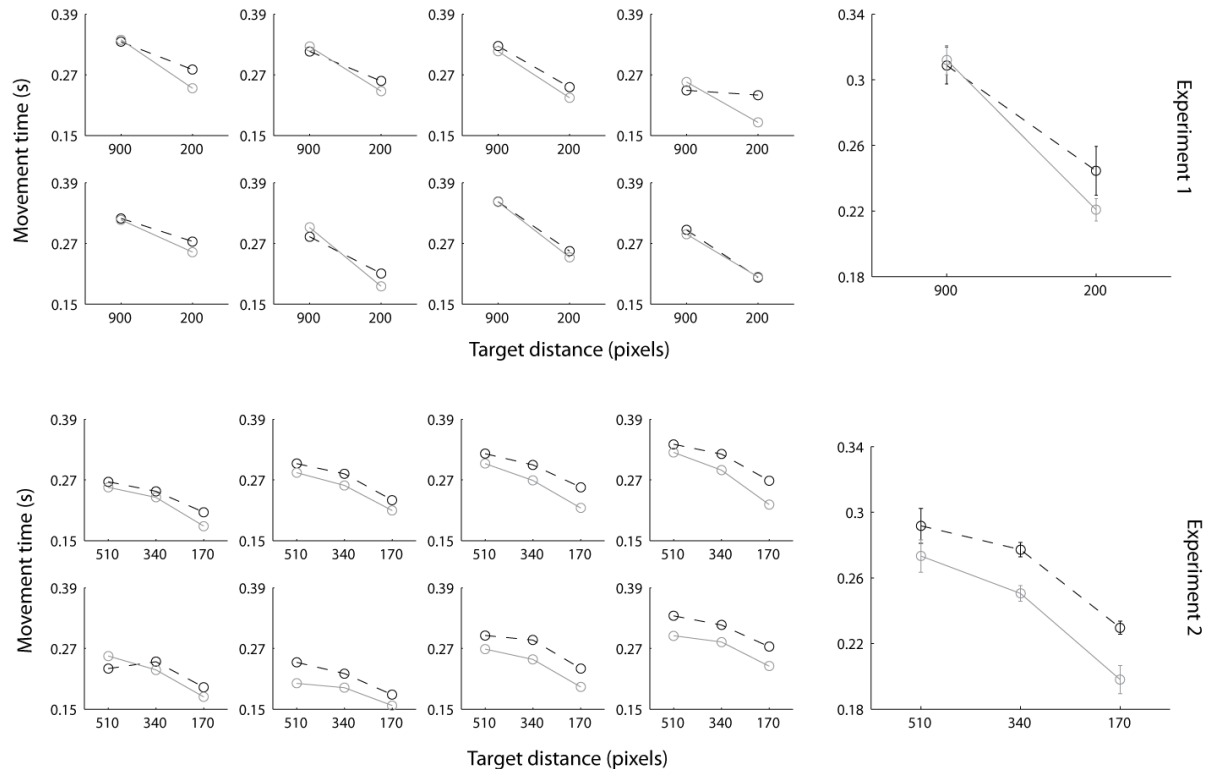


Fig S2.1. Movement time: group averages and individual movement time as a function of target distance, target size and experiment. The dashed line represents small targets and the full line represents large targets. Error bars are 95% confidence intervals useful for within-subject comparisons.

8.1.1.1 Movement time. As can be seen in Fig. S2.1, movement times show approximately the same pattern as the response times reported on in Chapter 2. When participants point to far targets they use more time than when pointing to near targets (Experiment 1: $F(1,7) = 120.68$, $p < .001$, $\eta_p^2 = .95$, Experiment 2: $F(2, 14) = 289.28$, $p < .001$, $\eta_p^2 = .98$). Recall that the time available for each pointing movement was the same regardless of distance. If participants had used nearly all the available response time (550 ms) to point at targets, the plots in Fig. S2.1 would have been horizontal lines. Faster movement times for near targets suggest that movements to near targets were faster than necessary. Another interesting trend is that participants moved more slowly

towards small targets than they did to large targets (Experiment 1: $F(1, 7) = 21.31$, $p = .002$, $\eta_p^2 = .75$, Experiment 2: $F(1,7) = 41.81$, $p < .001$, $\eta_p^2 = .86$). Given the speed-accuracy trade-off (Fitts, 1954; Schmidt et al., 1979) this also suggests that people were sacrificing precision for movement speed. In Experiment 1 there was also a significant interaction $F(1, 7) = 9.54$, $p = .018$, $\eta_p^2 = .58$), whereas in Experiment 2 it did not reach significance ($F(2, 14) = 41.81$, $p = .09$, $\eta_p^2 = .29$).




Fig S2.2. Reaction time: group averages and individual reaction time as a function of target distance, target size and experiment. The dashed line represents small targets and the full line represents large targets. Error bars are 95% confidence intervals useful for within-subject comparisons.

8.1.1.2 Reaction time. Reaction times for most participants show the opposite, albeit weaker, trend to that of movement time (Fig. S2.2). Firstly, reaction time decreases with increases in target distance (Experiment 1: $F(1,7) = 11.32$, $p = .012$, $\eta_p^2 = .62$, Experiment 2: $F(2, 14) = 20.01$, $p < .001$, $\eta_p^2 = .74$). Thus, participants to some extent appear to trade off movement time with response time, initiating movements faster to targets that are far away. Secondly, for some participants reaction times to small targets are faster than reaction times to large targets. This trend was significant in

Experiment 2 ($F(1, 7) = 8.125, p = .025, \eta_p^2 = .54$) and marginal in Experiment 1 ($F(1, 7) = 4.92, p = .06, \eta_p^2 = .41$). In Experiment 1 there was also a significant interaction ($F(1,7) = 5.72, p = .048, \eta_p^2 = .45$), whereas none was detected in Experiment 2 ($F(2,14) = .376, p = .693, \eta_p^2 = .05$).

8.1.2 Control experiment

Here we give brief details on the small control study reported on in the General discussion. Only substantial deviations from Experiment 1 and 2 are noted.

8.1.2.1 Participants & instructions. Five participants took part. All participants (except Participant 2 who was the author) were paid at an hourly rate of £10.

Participants were informed that speeded motor movements are variable and that even if one aims for the same spot on each reach the actual end point will deviate randomly from trial to trial. They were further informed about the difference between accuracy and precision. Participants were instructed to minimize the distance between each end point and the perceived target centre. All participants (except Participant 2) were naive as to the purpose of the study.

8.1.2.2 Stimuli, Experimental Design and Procedure. On each trial, one target disc was displayed to the left of the dock at a distance of ~9.2 cm (340 pixels) at one of five angles ($-15^\circ, -7.5^\circ, 0^\circ, +7.5^\circ, +15^\circ$). Targets were either small (radius ~2.9mm / 11 pixels) or large (radius ~5.9 mm / 22 pixels) yellow discs. On each trial a random target size and target angle was selected for presentation. In total, 300 small and 300 large non-late and non-anticipatory trials were collected.

Participants received feedback identical to that in Experiment 1 and 2 on where they hit the screen (but did not receive any points for hitting the targets as we wanted to minimize the incentive for satisficing).

8.1.2.3 Data analysis & Results. For each participant we collapsed across target angle, creating one small target and one large target distribution (see Gordon et al., 1994). On a group level, there was a small but detectable effect of target size on movement time ($t(4) = 4.18, p = .014$, mean difference = 4.2 ms), indicating that movements to smaller targets were slower than movements to large targets. We did not detect an effect of target size on either movement time or reaction time ($t(4) = 1.02, p = .37$, mean difference = 1.6 ms; $t(4) = 2.51, p = .066$, mean difference = 5.9 ms).

We compared each participant's movement variability for small targets to their movement variability for large targets using un-paired t-tests. The t-statistic was used to

derive JZS Bayes Factors (Rouder et al., 2009), which allow inferences in favour of the null as well as in favour of the alternative hypothesis.

Three of five participants reached with equal precision to small and large targets (JZS Bayes Factors > 3) and the evidence for two of five participants was inconclusive. If one performs the same analyses on the data for Experiment 1 and 2, the results are markedly different – most participants reached with greater precision to small targets (12 of 16, JZS Bayes factors < 0.33) and only three of twelve participants reach with equal precision to small and large targets (JZS Bayes Factor > 3). A group-level analysis provides similar evidence, showing that participants had a lower average difference (between small and large targets) in movement variability compared to those in Experiment 1 and 2 ($t(19) = 3.86$, $p = .001$, mean difference = 1.33)²².

²² A reviewer questioned whether precision differences may be affected by target distance (in Experiment 3 one mid-distance was used, whereas Experiment 1 and 2 used two and three different target distances respectively). As a control, we therefore fit bivariate Gaussians to the mid-distance data in Experiment 2 (the same distance as used here). The parameters of these maximum-likelihood fits were used to simulate participants in Experiment 2 reaching, the same number of times as here, to mid-distance targets only. Even when distance and sample size has been equated, the average precision difference between small and large targets is larger in Experiment 2 than here ($t(11) = 3.74$, $p = .003$, mean difference = .96).

8.2 Chapter 4 – Supplementary materials

8.2.1 SI Methods: Fitting the Weibull function to accuracy data

In fitting the Weibull function we generally follow Wichmann and Hill (2001) as described below. To model accuracy, as a function of time spent on a given task, we use the generic psychophysical model:

$$\varphi(x; \boldsymbol{\theta}) = \varphi(x; \alpha, \beta, \gamma, \lambda) = \gamma + (1 - \gamma - \lambda)F(x; \alpha, \beta) \quad \text{SI Eq. 4.1}$$

Here, the model $\varphi(x; \boldsymbol{\theta})$ specifies the relationship between the probability of responding correctly, p , and the response time imposed x . The shape of the underlying function is determined by the parameters α , β , γ , λ and function F . The third and fourth parameters (γ , λ) determine the lower and the upper bound respectively. Here $\gamma = .5$. Parameter λ was free but constrained (see below for details). This corresponds to a flat Bayesian prior on λ in the constrained range. For F we choose the Weibull function:

$$F(x; \alpha, \beta) = 1 - \exp\left[-\left(\frac{x}{\alpha}\right)^\beta\right], \quad 0 \leq x \leq \infty \quad \text{SI Eq. 4.2}$$

We fit the model, as described by the parameter vector $\boldsymbol{\theta} = (\alpha, \beta, \gamma, \lambda)$, to individual data using a maximum likelihood approach. We define three vectors, \mathbf{n} , \mathbf{y} , and \mathbf{x} , each of length K (the number of blocks of data collected in the assessment phase of the experimental session). Vector \mathbf{n} describes the number of trials in each block. The data are described in the vector \mathbf{y} which records the proportion of correct responses in each of the K blocks of data. Vector \mathbf{x} is defined as the response time (deadline) which was fixed within each block but varied across blocks. The likelihood of the data given the model is then defined as

$$L(\boldsymbol{\theta}, \mathbf{y}) = \prod_{i=1}^K p(y_i | \boldsymbol{\theta}) = \prod_{i=1}^K \left[\binom{n_i}{y_i n_i} \varphi(x_i; \boldsymbol{\theta})^{y_i n_i} [1 - \varphi(x_i; \boldsymbol{\theta})]^{(1 - y_i) n_i} \right] \quad \text{SI Eq. 4.3}$$

Maximising the likelihood, L , is equivalent to minimising the quantity l such that:

$$l(\boldsymbol{\theta}, \mathbf{y}) = -\log L(\boldsymbol{\theta}, \mathbf{y}) = \sum_{i=1}^K \left[\log \binom{n_i}{y_i n_i} + y_i n_i \log \varphi(x_i; \boldsymbol{\theta}) + (1 - y_i) n_i [1 - \varphi(x_i; \boldsymbol{\theta})] \right]$$

SI Eq. 4.4

SI Eq. 4.4 was minimised using a global optimisation algorithm (GlobalSearch, Matlab). The first three parameters in the model were constrained as follows: $\alpha_{\min} = .01$, $\alpha_{\max} = 5000$, $\beta_{\min} = .01$, $\beta_{\max} = 500$, $\gamma = .5$. For most stimuli asymptotic performance near 100% correct was likely. In such cases, asymptotic performance was constrained to lie in the 95-100% correct range ($\lambda_{\min} = 0$, $\lambda_{\max} = .05$) (Wichmann & Hill, 2001). For difficult tasks, in which asymptotic performance is likely to lie outside this range, we separately estimated asymptotic performance (120 trials under accuracy maximization instructions). This data was used to obtain a maximum likelihood estimate (MLE) of the 95% confidence interval of asymptotic performance, which was used as lower and upper limit for λ for these tasks, whenever the achieved accuracy for the longest imposed response time was lower than .95. Due to missing data, Participant 2 in Experiment 4 is an exception to this. However, this participant reached a high accuracy even for the hard tasks (standard intervals [0-.05] for λ were applicable).

Allowing asymptotic values other than full accuracy generally produced satisfactory results. However, for Participant 1, in Exp. 1 there were a few local minima for the hard task. The global minimum essentially resulted in a step-function (very large β), whereas a local minimum resulted in slopes similar to other observers' slopes. We elected to present the latter. Note - the qualitative result (that this participant deviates from optimality) does not change if the global-minimum function is used. Participant 6's function for the easy motion discrimination task in Exp 4 was also step-like. This fit, however, was stable and therefore presented as found.

8.2.2 SI Methods: Assessing goodness of fit

Deviance

$$D = 2 \sum_{i=1}^K \left\{ n_i y_i \log \left(\frac{y_i}{p_i} \right) + n_i (1 - y_i) \log \left(\frac{1 - y_i}{1 - p_i} \right) \right\} \quad \text{SI Eq. 4.5}$$

was used as a metric of goodness of fit (Wichmann & Hill, 2001). For each function that was fit (N=61), an empirical deviance score was computed using the recorded data and the best fit function (derived using methods outlined "Fitting the Weibull function to

accuracy data” above). This empirical deviance score was compared to a distribution of simulated deviance scores. For each data set, the best fit function was used as a generating function for new data sets ($N = 10\,000$). For each simulated data set, a function was fit and deviance was computed. This procedure results in a distribution of deviances, against which the empirical deviance can be compared. If the empirical deviance lay outside the 97.5th percentile of this distribution data sets were classed as unlikely to have been generated by the best fit function.

Fig. S4 M1 shows the critical deviance of these distributions as a function of the empirical deviance, for all functions that were fit. Data points that lie above the identity line are data sets for which we failed to reject the hypothesis that the data sets were not generated by the best fit function. Data points that lie below the identity line suggest that the best fit function is an unlikely generator of the data. As can be seen, the majority (56 of 61) of best fit functions were not classed as unlikely to have generated the data sets.

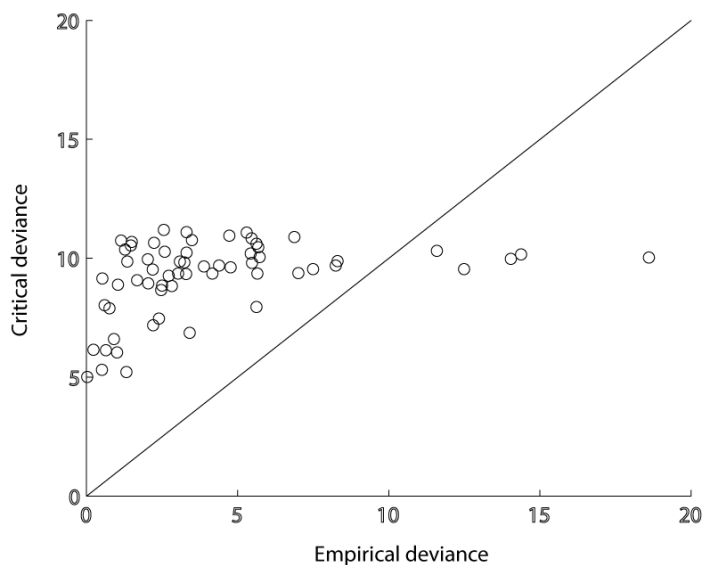


Fig. S4 M1. Model fits. 97.5 percentiles of the simulated deviance distributions as a function of the empirical deviances. Points above the identity line represent data sets for which we failed to reject the alternative hypothesis of a poor fit.

8.2.3 *SI Methods: Additional details*

8.2.3.1 Apparatus. Stimuli were viewed binocularly on a CRT with a red filter (1024 pixels x 768 pixels, 100 Hz) in a dark room (viewing distance 70 cm). Custom software (C++, OpenGL) displayed the stimuli, which was drawn in red. Participants responded by pressing buttons on a trackball.

8.2.3.2 Stimuli types & Tasks: direction discrimination. Random-dot patterns (see Shadlen & Newsom, 2001 for the algorithm) composed of dots were drawn on a black background within a circular aperture (visual angle = 19°), with a fixation dot (of different luminance). On each trial, the coherently moving dots moved (2 pixels/frame) left or right (with equal probability). Participants made left/right judgments (left button mapped to left movement). Dot coherences (proportion of dots moving coherently in one direction) varied across experiments and conditions.

In Exp. 3 the dot coherence was .3. We noticed large individual differences; with Participant 2 not reaching ~100% accuracy (we increased coherence for this participant to .5). For this reason, we tweaked the stimuli by increasing the dot density (Exp. 1-2 & 4: 160 dots/frame; Exp. 3: 80dots/frame). In experiments in which task difficulty was manipulated (Exp 1 and 4), dot coherence for the easy condition was .7 and dot coherence for the hard condition .2 (see SI Methods: Fitting the Weibull function to accuracy data for how the asymptote was allowed to vary beyond the standard 5% error rate for the latter condition). In Experiment 2 where difficulty was not manipulated dot coherence was .25.

8.2.3.3 Stimuli types & Tasks: mental arithmetic. Two numbers were presented centrally on each trial (~8° visual angle). Numbers were sampled from a uniform distribution (range 1 – 99), conditional on the sum of the numbers not being 100 and the absolute difference between the sum and 100, not being greater than 10. Participants judged whether the sum of the two numbers was smaller (left button) or larger (right button) than 100. In Exp. 4 the easy condition was created by sampling with a more limited precision (i.e., in 5's rather than 1's, e.g., [65, 30] vs. [66, 33]).

8.2.3.4 Stimuli types & Tasks: mental rotation. Two three-dimensional objects, composed of 10 cubes joined together (standard mental rotation figures, see Shepard & Metzler, 1967), were presented side by side (perspective projection). After both objects were given the same random orientations, one object was rotated by an additional 60° along its vertical axis. On each trial, objects were made impossible to align with a 50% chance.

8.2.3.5 Procedure. As sessions were completed on different days, experimental sessions began with some warm-up trials (10-30 depending on experiment) with no time limit, no rewards/penalties, and with auditory feedback to remind participants of the stimulus-response mappings. The N for each deadline in the assessment staged ranged from 60 to 100 (opportunistically) across experiments. In Experiment 3, responses that differed by ± 150 ms relative to deadline were re-run. As forced response times

increased, participants found it increasingly difficult to meet the same tight deadline. The tendency for increased variability with increased response times has been documented previously (Mates, Radil, Mueller & Poeppel, 1994). To avoid frustrating participants the deadline criteria for slower forced response times were relaxed in later experiments. Specifically: in Experiment 1, 2 and 4 the deadline limits were as follows: $\text{deadline} < 1000 \text{ ms} = \pm 100 \text{ ms}$, $1000 < \text{deadline} < 2500 = \pm 150 \text{ ms}$ and $\text{deadline} > 2500 = \pm 200 \text{ ms}$.

8.2.4 SI Methods: Mathematical formulation of the decision problem

We express the decisions made by our participants as choices between infinitely many lotteries (one for each possible average response time) with two possible outcomes. For a given average response time, RT , the associated lottery, $L(RT)$, is defined as:

$$L(RT) = \{(o_1, p_1(o_1|RT)), (o_2, p_2(o_2|RT))\} \quad \text{SI Eq. 4.6}$$

The outcome o_1 , is the penalty associated with responding incorrectly, occurring with conditional probability $p_1(o_1|RT)$ and outcome o_2 is the reward associated with responding correctly, occurring with conditional probability $p_2(o_2|RT) = 1 - p_1(o_1|RT)$. The expected value $EV(RT)$ for a given response time RT can then be expressed as:

$$EV(RT) = \sum_{i=1}^2 p_i(o_i|RT) o_i \quad \text{SI Eq. 4.7}$$

The $p_i(o_i|RT)$ are determined from the best fit Weibull function to the data obtained in the assessment stage.

Given the inter-stimulus interval, I , and the total time available in which to complete tasks, T , for each observer we can calculate the average number of responses, N , made in the time available as:

$$N(RT) = \frac{T}{I+RT} \quad \text{SI Eq. 4.8}$$

We can then calculate the overall expected gain, G , of each participant as the product of N and EV .

$$G(RT) = N(RT) \cdot EV(RT) = \left(\frac{T}{I+RT} \right) \sum_{i=1}^2 p_i(o_i|RT) o_i \quad \text{SI Eq. 4.9}$$

We defined efficiency, $E(R)$ as a re-scaled expression of the overall expected gain over the range 100ms-5000ms, which is a reasonable time-span given the experimental parameters. The efficiency is computed as:

$$E(RT) = \frac{G(RT) - \min_{100 \leq R \leq 5000} G(RT)}{\max_{100 \leq R \leq 5000} G(RT) - \min_{100 \leq R \leq 5000} G(RT)} \quad \text{SI Eq. 4.10}$$

The optimal response time is the particular response time which maximizes efficiency over this range.

For Experiment 4, SI Eq. 4.9 was extended to take into account the dependence between easy (RT_1) and hard (RT_2) choices performed in the same period (T). Outcomes and conditional probabilities for RT_1 are as above, and RT_2 is also associated with a positive (o_3) and a negative (o_4) outcome and their respective conditional probabilities. SI Eq. 10 was extended similarly.

$$G(RT) = \left(\frac{T}{\sum_{j=1}^2 I_j + RT_j} \right) \sum_{i=1}^4 p_i(o_i|RT) o_i \quad \text{SI Eq. 4.11}$$

Due to the experimental design (unpredictable trial-by-trial changes in task difficulty together with self-paced trial display) we were unable to match the number of hard and easy trials a priori. Thus, SI Eq. 4.11 holds in terms of expectation only. Across observers in Experiment 4 the proportion of experienced easy trials was .5026 (total $N = 2505$), the maximum likelihood 95 percentile of this estimate is .48 – .52. The percentiles for estimates for individual observers similarly cover .5. Consequently, although the frequency of easy and hard trials was not precisely matched the counts were similar. Therefore, it does not seem inappropriate to define the expected gain using SI Eq. 4.11.

8.2.5 SI Methods: Extrapolation from Weibull fits to choice data

To assess whether there were systematic differences between predicted- and actual accuracy, we compared the predicted accuracy, for all response time choices, to

the actual accuracy. This comparison is shown in Fig. S4 M2 and was made across all experiments, tasks, conditions and participants – in total 121 accuracy and predicted accuracy pairs (this also includes data from a reward manipulation for the mental-arithmetic and rotation tasks, not reported on in the main paper, but reported in Fig. S4.5). As can be seen the distribution of differences (predicted-obtained) is captured by a Gaussian centred near 0 – indicating that predicted accuracies do not systematically deviate from obtained accuracies (and that the psychometric functions therefore provide unbiased estimates of accuracy as a function of response time).




Fig. S4 M2. Differences between predicted accuracy and obtained accuracy (Δ Accuracy) across all experiments, conditions and observers (N=121). The line is a Gaussian fit to the data (MLE), with parameters $\mu = -.002$ and $\sigma = .055$. The 95% confidence interval on μ was $-.012$ to $.008$ ($\sigma = .049$ to $.063$).

8.2.6 SI Methods: Was the hard-easy manipulation successful?

Visual comparison of hard and easy accuracy functions suggested that all task difficulty manipulations were successful (e.g., see Fig. 2A). We tested whether the apparent differences were significant with a Monte Carlo hypothesis test under the null hypothesis that the hard and easy data sets were generated by a common underlying function (MLE fit to the combined hard-easy data set). That is, the empirical difference between the best-fit parameters for the hard- and easy task was evaluated against a null distribution of parameter differences, under the assumption that both the hard- and easy data sets were generated by the same underlying function.

Briefly, hard- and easy data sets were combined into joint data sets. Each combined data set was fit, and the best fit functions were used to generate $2 \times 10\,000$ data sets (one vector for easy- and one for hard tasks). Each of these data sets was fit anew. The difference between the resulting best fit parameter vectors ($D(\alpha_\Delta, \beta_\Delta, \lambda_\Delta) = E(\alpha_{\text{easy}}, \beta_{\text{easy}}, \lambda_{\text{easy}}) - H(\alpha_{\text{hard}}, \beta_{\text{hard}}, \lambda_{\text{hard}})$) is the null-distribution. The differences between the empirical hard- and easy parameter vectors, of functions fit to the original hard- and easy data sets, were then evaluated against a three-dimensional kernel density estimate of this null-distribution. If the empirical difference was unlikely given the null-distribution ($p < .05$) then the two functions were classed as different. By this criterion, 15 of 18 hard-easy accuracy function-pairs had statistically different parameters.

8.2.7 Fig. S4.1.

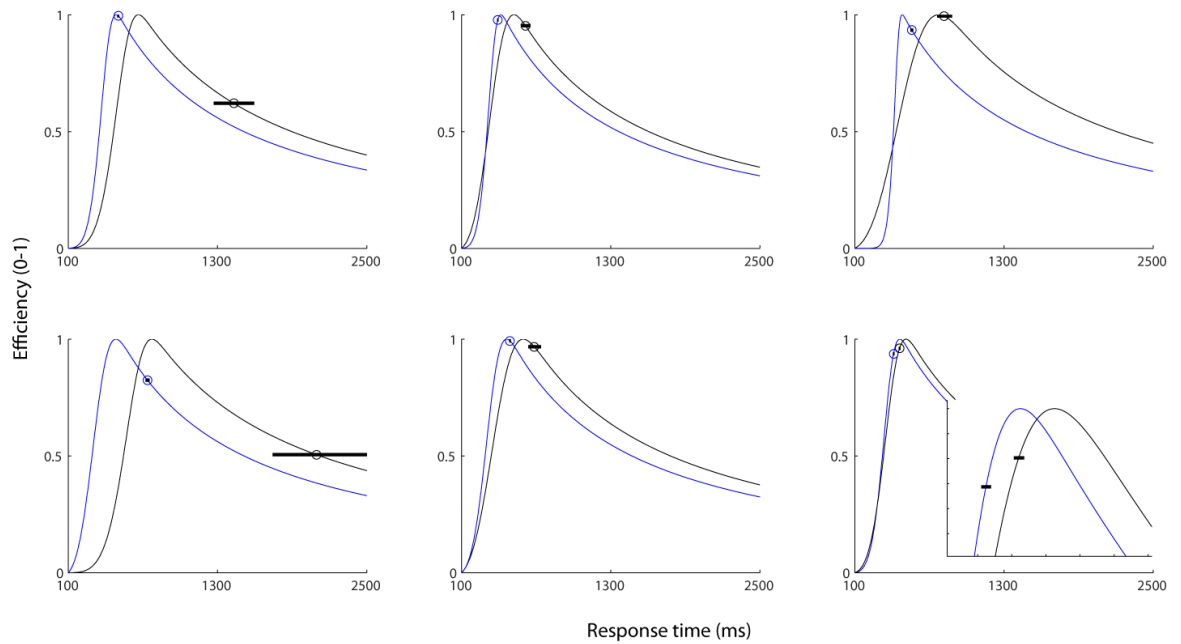


Fig. S4.1 Efficiency functions for data reported in Fig 2C. Efficiency function for easy- (blue) and hard (black) motion discrimination, with average response times (circles) including bootstrapped 95% CI's, for participant 1-6 (A-F). For participant 6, the inset shows a zoomed in image of the peaks. As can be seen, participants always shifted in the right direction for the hard blocks relative to the easy blocks (i.e., slowed down). For most participants, responses were also near the peak of the functions. We confirmed that choices (circles) were closer to the appropriate efficiency function peak than they were to the wrong efficiency function peak ($t(5) = -3.21, p = .024$). Similarly, a simple linear model (with two multivariate outliers deleted: Mahalanobis distance Chi-Square criterion of $p > .01$), of choices and peaks which assumes that participants base their choices on the correct efficiency function peaks, fits the data well ($r(8) = .78, p = .008$). A model assuming that participants base their choices on the wrong efficiency function peaks, on the other hand, fits the data less well ($r(8) = .15, p = .69$).

8.2.8 Fig. S4.2




Fig. S4.2. Response time choices for the initial reward structure manipulation. White symbols correspond to timing choices for the three conditions. Gray triangles correspond the optimal response time for each condition and participant. The distance of the actual choice (white symbols) to the gray triangle is a measure of how far away from optimal participants choices are. All error bars are bootstrapped 95% CI's. Some error bars are too small to be visible at this scale.

8.2.9 Fig. S4.3




Fig. S4.3. Efficiency functions for data reported in Fig S4.2. Efficiency functions for participant 1 – 5 (rows), for the neutral (black), the penalty only (blue) and the no-penalty (red) condition, with average response times (circles). Error bars are bootstrapped 95% CI's (barely visible at this scale)

8.2.10 Fig. S4.4

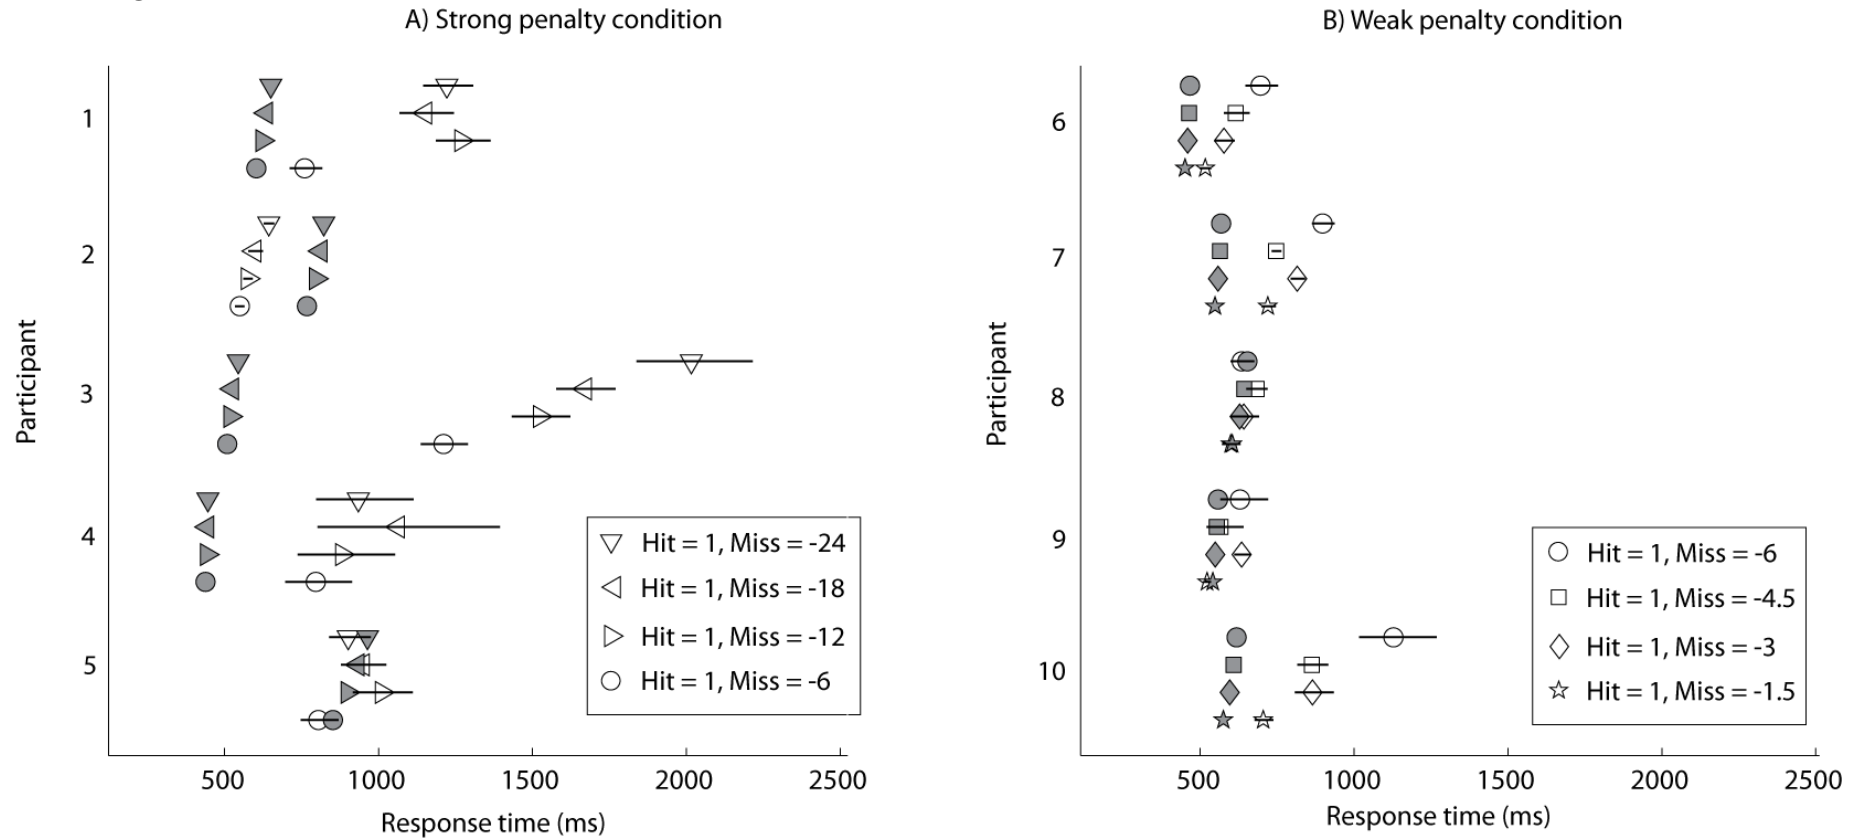


Fig S4.4. Response time choices for the between-subject penalty magnitude manipulation. White symbols correspond to timing choices for the four penalty levels. Gray symbols correspond the optimal choices for each condition and participant. The distance of the actual choice (white symbols) to the gray symbols is a measure of how far away from optimal participants choices are. All error bars are bootstrapped 95% CI's.

8.2.11 Fig. S4.5

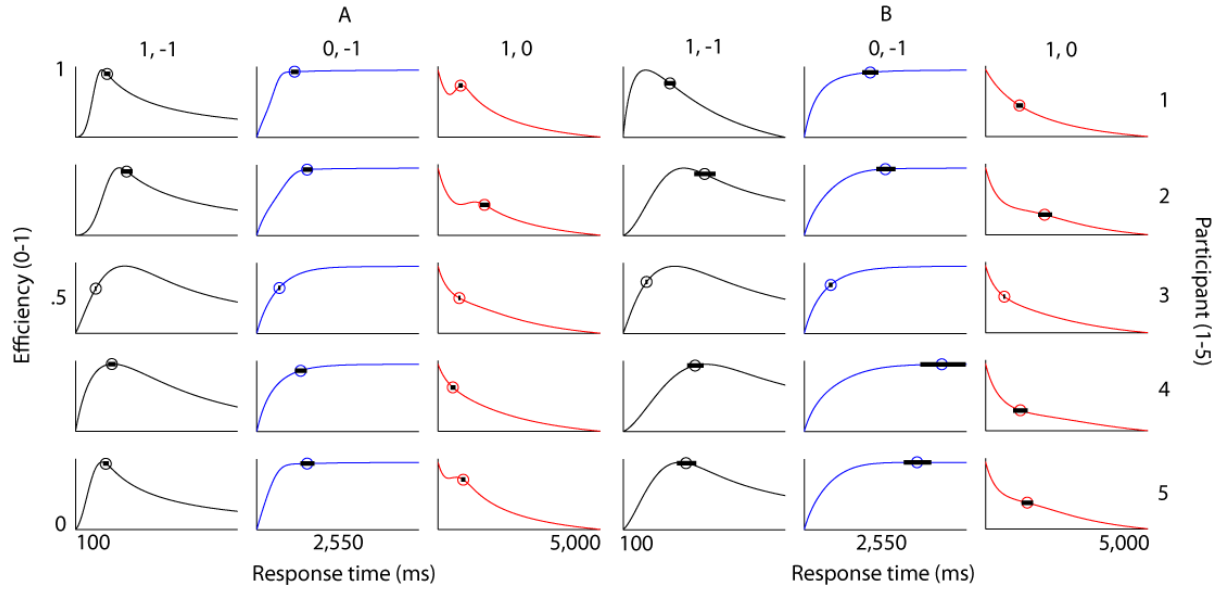


Fig. S4.5. Cognitive task efficiency functions for the pilot reward structure manipulation. Mental arithmetic (A) and mental rotation (B), for participant 1 – 5 (rows), for each condition: neutral (black), penalty-only (blue) and reward-only (red). As can be seen, the effect of choosing sub-optimally is minor for the penalty-only condition. Participants shift their response times sufficiently far given the flatness of the efficiency functions (except Participant 3, who is consistently aggressive across each task and reward structure [as they were for the motion discrimination task, see Fig. S4.4]). However, the sub-optimal response time shifts for the reward-only task (red) now become consequential: with many participants now earning less than 50% of the maximum earnings (i.e., the optimal response for the reward-only condition has now become as extreme as for the penalty-only condition). This is due to the relative flatness of the underlying time-accuracy functions (compared to motion discrimination) and due to, relative to the asymptote of the time-accuracy function (and hence relative to the motion discrimination task), the short inter-stimulus interval (for a related point see, Green, 1960; Winterfeldt & Edwards, 1968).

8.3 Chapter 6 – Supplementary materials

8.3.1 SI Additional methods: Model fitting

As noted in Chapter 6, the space of possible cumulative prospect theory (Tversky & Kahneman, 1992) parameterizations is large. It is arguably not feasible to search the entire space. Instead we picked three model parameterizations: the original parameterization (Tversky & Kahneman, 1992) with a choice function, the parameterization recommended by Stott (2006, the one reported in Chapter 6) and Wu et al.'s (2009) parameterization. To ensure that the effect of particular aspects of each parameterization was balanced we fit a full factorial combination of the parameterizations of these three models. As outlined in Table S1 below, this procedure resulted in 8 different model parameterizations.

Table S1. Illustration of the combinations of cumulative prospect theory parameterizations that was fit.

<i>Nr</i>	<i>Value function</i>	<i>Probability function</i>	<i>Choice Function</i>	<i>Noise type</i>
1	Power	Prelec	Wu et al.	Proportional
2	Power	Prelec	Logistic	Proportional
3	Power	Prelec	Wu et al.	Constant
4	Power	Prelec	Logistic	Constant
5	Power	Kahneman & Tversky	Wu et al.	Proportional
6	Power	Kahneman & Tversky	Logistic	Proportional
7	Power	Kahneman & Tversky	Wu et al.	Constant
8	Power	Kahneman & Tversky	Logistic	Constant

In the following we will use x and p to denote values and probabilities respectively. The three target model parameterizations all use the same power function Tversky & Kahneman (1992) use. Because we used only non-negative values ($x \geq 0$) we estimate α (and not λ or β):

$$v(x) = \begin{cases} x^\alpha & \text{if } x \geq 0 \\ -\lambda(-x)^\beta & \text{if } x < 0 \end{cases} \quad \text{SI Eq. 1}$$

The original probability weighting function considered by Tversky & Kahneman (1992) is:

$$w(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{1/\gamma}}, \quad \text{SI Eq. 2}$$

For $\gamma > 1$ low probabilities are underweighted and for $\gamma < 1$ low probabilities are overweighted.

Prelec's (1998) one-parameter probability weighting function, recommended by Stott (2006) and fit by Wu et al. (2009) is:

$$w(p) = e^{(-(-\ln p)^\gamma)} \quad \text{SI Eq. 3}$$

As above, $\gamma > 1$ implies underweighting and $\gamma < 1$ implies overweighting.

A choice function maps properties of prospects onto choice probabilities. Stott (2006) recommends a logistic choice function. The probability of choosing prospect B, when faced with the choice between it and prospect A, for the logistic function is:

$$p_B = \frac{1}{1 + e^{k(\varphi(A) - \varphi(B))}}, \quad \text{SI Eq. 4}$$

The probability of choosing prospect A is $1 - P_B$. Parameter k in this function can be thought of as a noise parameter. The lower k is the worse we become at discriminating between the two prospects. This way of modelling choice is typically combined with the assumption that the noise (k) is constant across all prospect pairs.

Wu et al. (2009) instead model noise as proportional to the prospects. They introduced proportional error by modelling each prospect as a random variable with a variance dependent on the prospect and a constant k . For example, prospect B can be expressed as $\psi(B) = N(v(x_B)w(p_B), k v(x_B)w(p_B))$. The difference between two such prospects, or random variables, is another random variable ($\Delta = \psi(A) - \psi(B)$), with variance equal to the sum of the variance of the two independent variables $\sigma_\Delta^2 = \sigma_B^2 + \sigma_A^2$. The probability of choosing prospect B then is the integral of Δ from 0 to negative infinity:

$$p_B = \int_{-\infty}^0 \frac{1}{\sqrt{2\pi\sigma_\Delta^2}} e^{-\frac{(\delta-\Delta)^2}{2\sigma_\Delta^2}} d\delta, \quad \text{SI Eq. 5}$$

The probability of choosing prospect A is $1 - P_B$.

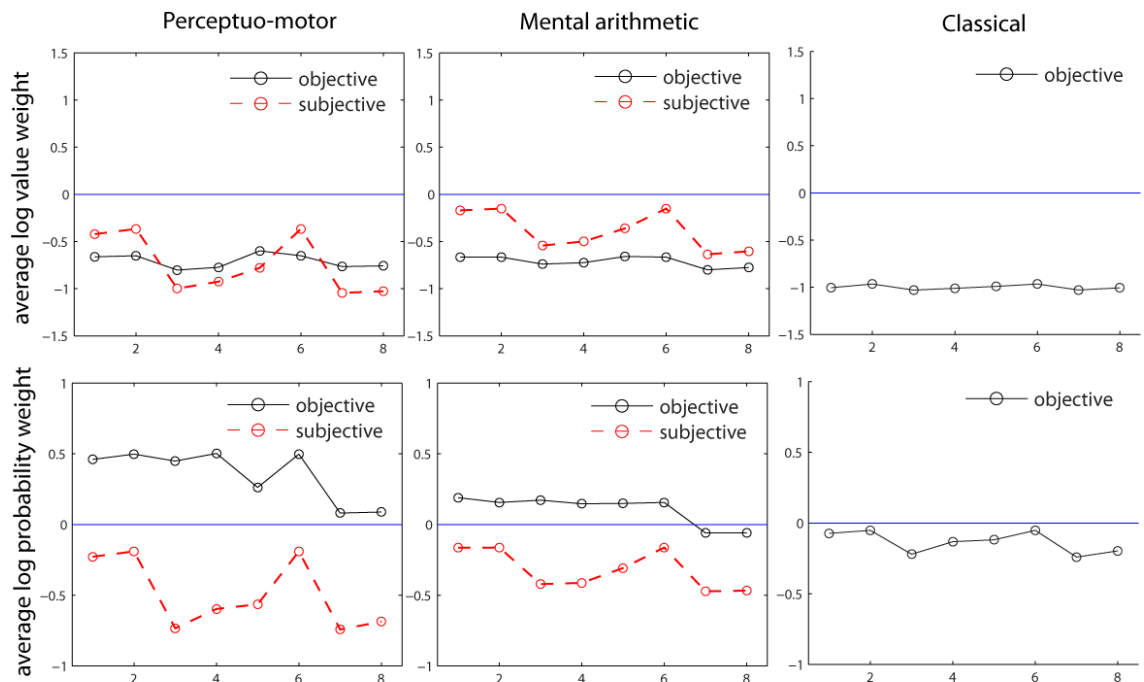
Thus Stott and Wu et al. recommend/use different choice functions and model noise in different ways. To model constant noise for Wu et al.'s choice function we made the variance of the choice variable independent of the prospects (e.g., $\psi(B) = N(v(x_B)w(p_B), k)$). To model proportional noise for the logistic choice function we modelled the prospects as random variables (as Wu et al. did) and made the parameter k equal to the precision ($1/\sigma^2_\Delta$) of the choice variable. These parameterizations were combined as noted in Table S1 above.

We fit the above eight parameterizations to individual participants' choices separately for each task by maximum likelihood methods using Matlab's Multistart solver (using 3000 solvers). As Wu et al. (2009), we denote a choice of prospect A as r and a choice of prospect B as $1-r$. We minimized the negative log-likelihood of the value weight (α), the probability weight (γ) and the noise weight (k) given participants choices:

$$-L(\alpha, \gamma, k) = - \sum_{i=1}^n (r_i \log(P_{A_i}) + (1 - r_i) \log(P_{B_i})) \quad \text{SI Eq. 6}$$

We constrained the probability and value parameters in the above functions to lie between .01 and 100; a range sufficient to capture both extreme under and overweighting. The proportional noise parameter was constrained to lie between 1e-10 and 100. The constant noise parameter was constrained to lie between 0 and 500 (for Wu et al.'s choice function the minimum was always 1e-10, whereas for the logistic function the minimum was 0). The difference between the ranges for the noise parameters is due to the proportionality of the proportional error model. Constraining the parameters as above is equivalent to applying a uniform Bayesian prior over the constrained ranges. As long as the constraints are reasonable, constraining the parameters will improve the speed at which global optima can be found with no ill effects.

Thus to ensure that the results reported on in the main paper were not due to the fitting of a specific parameterization of cumulative prospect theory, we fit the eight parameterizations just outlined. The results of this exercise are shown in Fig. SM6.1 below.



Cumulative prospect theory parameterization - see Table S1 for key

Fig. SM6.1 Best fit value and probability weights for eight different parameterizations of cumulative prospect theory. The Y-axes show the average (50% trimmed mean) of the logarithm of the value (Row 1) and the probability (Row 2) weights. A log weight of 0 implies that participants were on average unbiased (had weights of 1). A log probability weight below 0 implies overweighting of low probabilities (as typical in classical tasks) and a log probability weight above 0 implies underweighting of probabilities (as typical in decisions from experience). The logarithm of the parameter is appropriate as it emphasizes under- and overweighting to an equal degree. The X-axis maps onto specific cumulative prospect theory parameterizations (see Table S1 for a key). Black full lines show average best-fit weights to objective probabilities. Red dashed lines show average best-fit weights to subjective probabilities.

As can be seen (Fig. SM6.1), nearly all parameterizations show the same qualitative pattern: When objective probabilities are modelled (black discs, full lines), average probability weights for decisions based on numerical probabilities (classical) and those based on internal estimates of probabilities (perceptuo-motor and arithmetic) dissociate. For classical decisions average log probability weights are negative but for arithmetic and for perceptuo-motor decisions the average weights are positive.

However, when subjective probabilities are fit (red discs, dashed lines) average probability weights show the same qualitative pattern across all three decision types. That is, when subjective probabilities are modelled the average participant overweights

small probabilities - regardless of task (i.e., has negative log-probability weights). Thus, across eight cumulative prospect theory parameterizations the results in Fig. 6.3 and Fig. 6.5 replicate.

8.3.2 Fig. S6.1

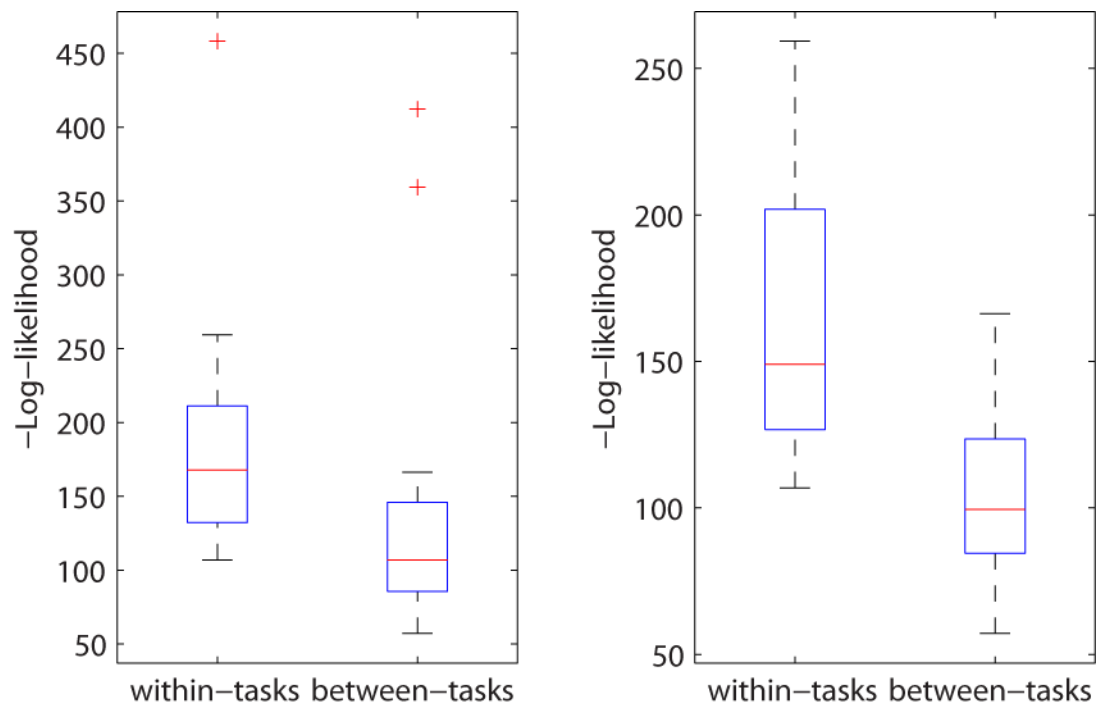


Fig. S6.1 Negative log-likelihoods for the within-between contrast. Left panel – full data set with three outliers identified (crosses are more than 2 inter-quartile ranges from the median). Right panel – data with the three outliers removed. The left box plot in each panel shows the negative log-likelihoods for the within-task-across-participants-predictions and the right box plot in each panel shows the across-tasks-within-participants-predictions.

8.3.3 Fig. S6.2




Fig. S6.2. Illustrations of perfect group-calibration with poor individual calibration. Three examples (one in each panel) with two participants who either show marked biases (Panel 1 & 3) or are very noisy (Panel 2) and yet give rise to perfect group-calibration. That is, once the two lines have been averaged, the average would coincide with the identity line. Note, scales are arbitrary and lines were hand-drawn.

8.3.4 Fig. S6.3

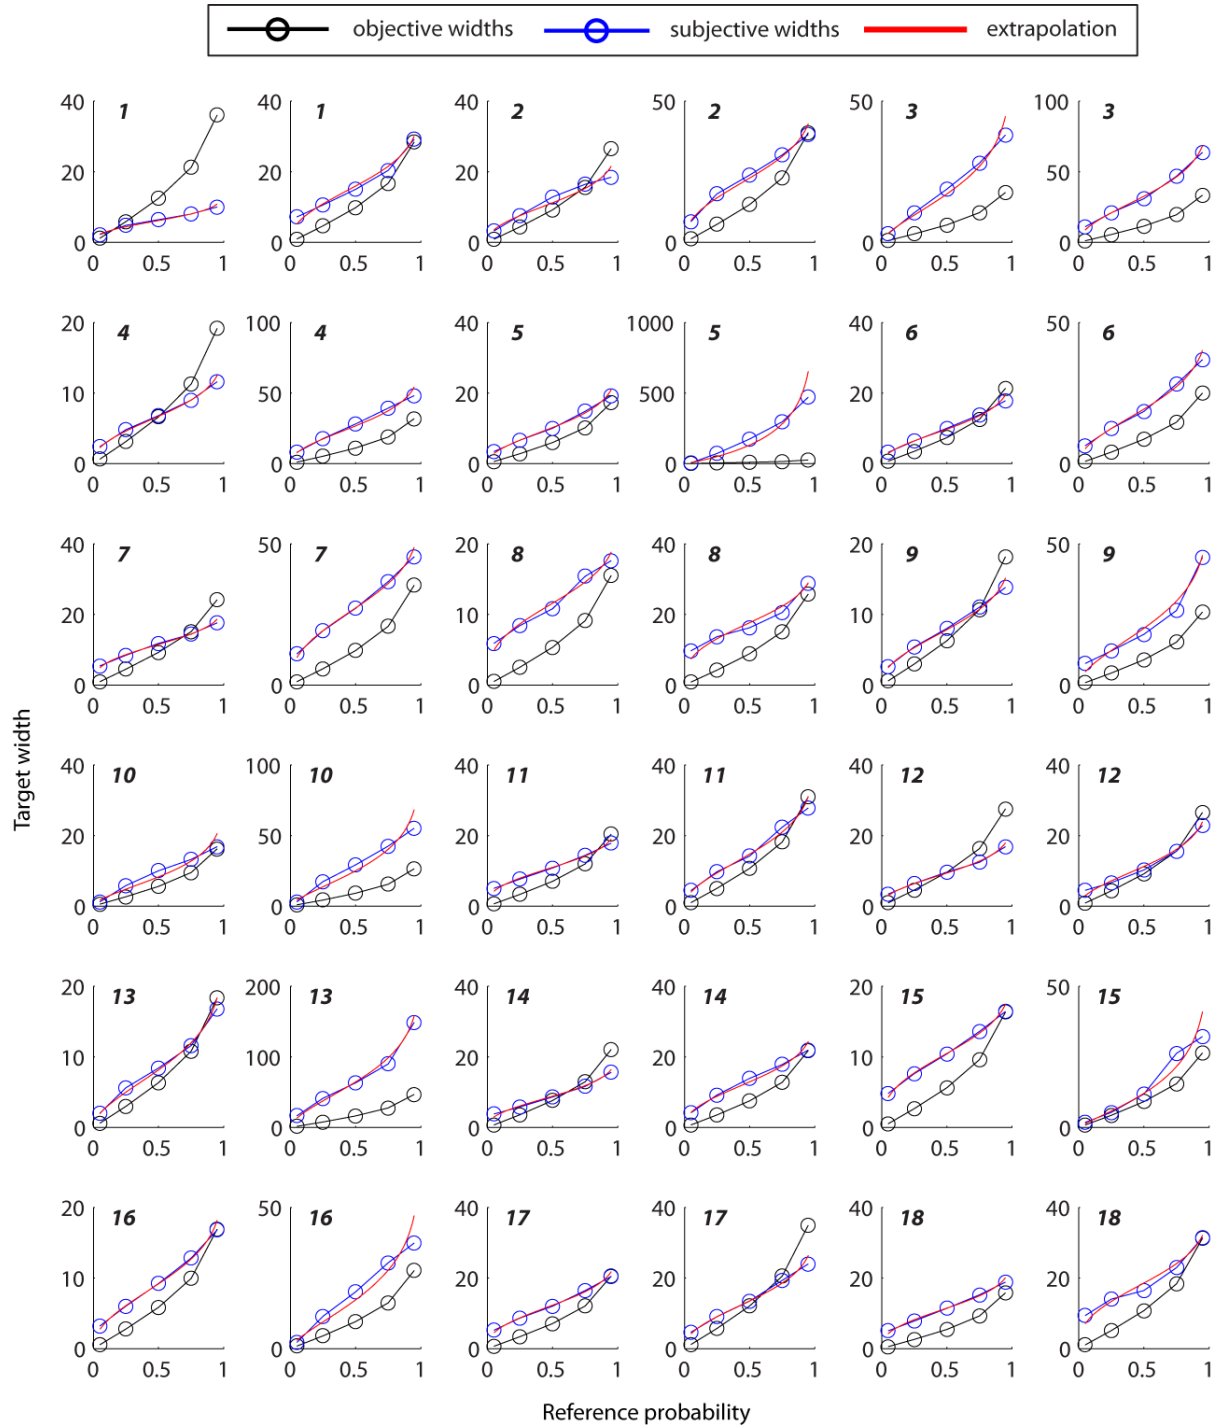


Fig. S6.3. Objective and subjective target widths as a function of hit probability for individual participants. The first panel in each pair shows calibration for the mental arithmetic task. The second panel in each pair shows calibration for the pointing task. Bold numbers correspond to participants. Black circles represent objective widths. That is, the target widths participants need to hit the targets with .05, .25, .5, .75 and .95 probability respectively. The blue circles represent the average of participants 5 last (of 6) width ratings. The red lines are extrapolations based on the Weibull function. Note, the y-axis scale differs from plot to plot.

8.3.5 Fig. S6.4




Fig. S6.4. Group-level calibration. Plots obtained by first averaging over each subjects 5 estimates for each of the 5 reference probabilities. This results in 18 estimates for each reference probability. These were mapped onto probability space and averaged to produce group-calibration curves.

8.3.6 Fig. S6.5

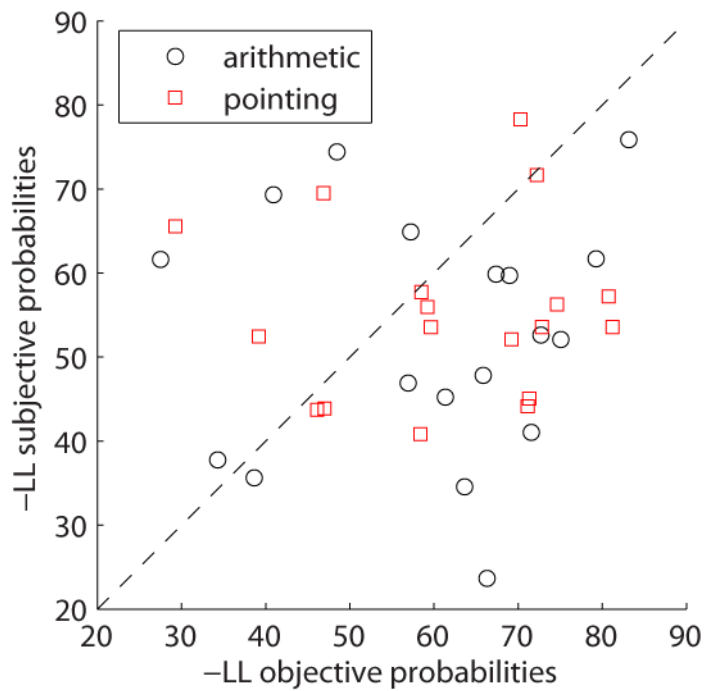


Fig. S6.5. Negative log-likelihoods for objective- and subjective probability fits. Arithmetic and pointing fits are colour and shape coded. Each symbol is a fit to one participant's data. The fit model is a simple expected value maximization model with proportional noise. As can be seen, most data points lie below the identity line indicating that probabilities in the form of our participants' beliefs about their ability to hit targets (subjective) better account for the data than probabilities in the form of their actual ability to hit targets (objective).

8.3.7 Fig. S6.6

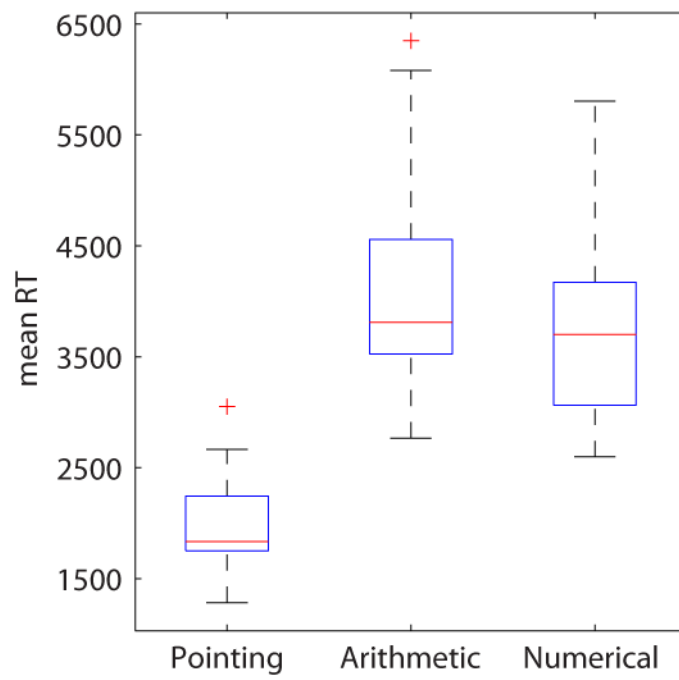


Fig. S6.6. Mean response times (ms) as a function of decision type.

9. References

- Abdellaoui, M., L'Haridon, O., & Paraschiv, C. (2011). Experienced vs. described uncertainty: Do we need two prospect theory specifications. *Management Science*, *57*, 1879–1895.
- Allais, M. (1952/1979). The foundations of a positive theory of choice involving risk and a criticism of the postulates and axioms of the American school. In M. Allais and O. Hagen (eds.), *Expected utility and the Allais Paradox*, 27 - 145. New York: Springer.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Arrow, K. J., (1973). Optimal insurance and generalized deductibles. R-1108-OEO, A report prepared for Office of Economic Opportunity, RAND Corporation: Santa Monica, CA.
- Barlow, H. B. (1962). Measurements of the quantum efficiency of discrimination in human scotopic vision. *Journal of Physiology*, *160*, 159–188.
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, *16*, 215–233.
- Battaglia, P. W., & Schrater, P. R. (2007). Humans trade off viewing time and movement duration to improve visuomotor accuracy in a fast reaching task. *The Journal of Neuroscience*, *27*, 6984–6994.
- Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis (2nd Ed)*. Springer-Verlag: New York.
- Bernoulli, D. (1738/1954). Exposition of a new theory on the measurement of risk. *Econometrica*, *22*, 23-36.
- Binstead, G., Brownell, K., Vorontsova, Z, Heath, M., & Saucier, D. (2007). Visuomotor system uses target features unavailable to conscious awareness. *Proceedings of the National Academy of Sciences*, *104*, 12669–12672.
- Birnbaum, M. H. (2008). New paradoxes of risky decision making. *Psychological Review*, *115*, 463–501.
- Bogacz, R. (2007). Optimal decision-making theories: Linking neurobiology with behaviour. *Trends in Cognitive Sciences*, *11*, 118–125.
- Bogacz, R., Hu, P. T., Holmes, P. J., & Cohen, J. D. (2010). Do humans produce the speed-accuracy trade-off that maximizes reward rate? *Quarterly Journal of*

- Experimental Psychology*, 63, 863–891.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 433–436.
- Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2008). Risky choice with heuristics: Reply to Birnbaum (2008), Johnson, Schulte-Mecklenbeck, and Willemsen (2008), and Rieger and Wang (2008). *Psychological Review*, 115, 281–290.
- Camerer, C. F., & Hogarth, R. M. (1999). The effect of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty*, 19, 7 - 42.
- Camilleri, A., & Newell, B. (2011). When and why rare events are underweighted: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psychonomic Bulletin & Review*, 18, 377– 384.
- Chater, N., & Oaksford, M. (2008). The probabilistic mind: where next? In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 501–514). Oxford: Oxford University Press.
- Chu, Y-P., Chu, R-L. (1990). The subsidence of preference reversals in simplified and marketlike experimental settings: A note. *The American Economic Review*, 80, 902–911.
- Churchland, M., Afshar, A., & Shenoy, K. (2006). A central source of movement variability. *Neuron*, 52, 1085–1096.
- Cohen, M. S., Kosslyn, S. M., Breiter, H. C., DiGirolamo, G. J., Thompson, W. L., Anderson, A. K., Bookheimer, S. Y., Rosen, B. R., & Belliveau, J. W. (1996). Changes in cortical activity during mental rotation: A mapping study using functional MRI. *Brain*, 119, 89–100.
- Dean, M., Wu, S-W., & Maloney, L. T. (2007). Trading off speed and accuracy in rapid goal-directed movements. *Journal of Vision*, 7, 1–12.
- Dickhaut, J., Rustichini, A., & Smith, V. (2009). A neuroeconomic theory of decision process. *Proceedings of the National Academy of Sciences USA*, 106, 22145–22150.
- Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. New York, NY: Chapman and Hall.
- Elliott, D., Hansen, S., Mendoza, J., & Tremblay, L. (2004). Learning to optimize speed, accuracy and energy expenditure: A framework for understanding speed-accuracy relations in goal-directed aiming. *Journal of Motor Behavior*, 36, 339–351.
- Elliott, D., Helsen, W. F., & Chua, R. (2001). A century later: Woodworth's (1989)

- two-component model of goal-directed aiming. *Psychological Bulletin*, 127, 342–357.
- Ellsberg, D. (1961). Risk, Ambiguity and the Savage Axioms. *Quarterly Journal of Economics*, 75, 643-669.
- Erev, I., Roth, A. E., Slonim, R. & Barron, G. (2002). Combining a theoretical prediction with experimental evidence. *Working paper*, downloaded from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1111712 on 02.01.2012
- Evans J. St. B. T. (1993). Bias and rationality. In K. I. Mankeltow., & D. E. Over (Eds.). *Rationality: Psychological and philosophical perspectives*. (pp. 6–30). London: Routedledge.
- Evans J. St. B. T., & Over, D. E. (1996). *Rationality and reasoning*. Hove, UK: Psychology Press.
- Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews: Neuroscience*, 9, 292–303.
- Feldman, R. S., (2006). *Understanding Psychology*, 8th Ed. McGraw-Hill: New York.
- Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movements. *Journal of Experimental Psychology*, 47, 381–391.
- Fitts, P. M., & Petersen, J. R. (1964). Information capacity of discrete motor responses. *Journal of Experimental Psychology*, 67, 103–112.
- Fox, C. R., & Hadar, L. (2006). "Decisions from experience" = sampling error + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making*, 1, 159–161.
- Fox, C. R., & Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management Science*, 44, 879-895.
- Gallistel, C. R. (2009). The importance of proving the null. *Psychological Review*, 116, 439–453.
- Geisler, W. S. (2003). Ideal observer analysis. In: L. Chalupa and J. Werner (Eds.), *The visual neurosciences* (825–837). Boston: The MIT Press
- Gepshtein, S., Seydell, A., & Trommershäuser, J. (2007). Optimality of human movement under natural variations of visual-motor uncertainty. *Journal of Vision*, 7, 1–18.
- Gigerenzer, G., (1996). On narrow norms and vague heuristics: A reply to Kahneman & Tversky (1996). *Psychological Review*, 103, 592-596.

- Gigerenzer, G., (2008). *Rationality for mortals: How people cope with uncertainty*. Oxford University Press: Oxford.
- Glimcher, P. W., Camerer, C. F., Fehr, E., & Poldrack, R. A. (2009). *Neuroeconomics: Decision making and the brain*. London: Academic Press.
- Gonzales, R., & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology*, 38, 129-166.
- Gordon, J., Ghilardi, M. F., & Ghez, C. (1994). Accuracy of planar reaching movements: Independence of direction and extent variability. *Experimental Brain Research*, 99, 97–111.
- Gottlieb, D. A., Weiss, T., & Chapman, G. B. (2007). The format in which uncertainty information is presented affects decision biases. *Psychological Science*, 18, 240-246.
- Green, D. (1960). Psychoacoustics and detection theory. *Journal of the Acoustical Society of America*, 32, 1189–1203.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: John Wiley and Sons.
- Gribble, P. L., Mullin, L. I., Cothros, N., & Mattar, A. (2003). Role of cocontraction in arm movement accuracy. *Journal of Neurophysiology*, 89, 2396–2405.
- Grice, G. R. (1968). Stimulus intensity and response evocation. *Psychological Review*, 75, 359–373.
- Harris, A. J. L., Corner, A., & Hahn, U. (2009). Estimating the probability of negative events. *Cognition*, 10, 51-64.
- Harris, C. M., & Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, 394, 780–784.
- Hau, R., & Pleskac, T. J., Kiefer, J., Hertwig, R. (2008). The description-experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, 21, 493-518.
- Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, 13, 517–523.
- Hertwig, R., & Ortmann, A. (2001). Experimental practises in economics: A methodological challenge for psychologists. *Behavioral and Brain Sciences*, 24, 383 - 451.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15, 534–539.

- Hilton, D.J. (1995). The social context of reasoning: Conversational inference and rational judgment. *Psychological Bulletin*, *118*, 248–271.
- Hudson, T. E., Maloney, L. T. & Landy, M. S. (2008). Optimal movement timing with temporally asymmetric penalties and rewards. *PLoS Computational Biology*, *4*, 1–9.
- Jensen, N. E. (1967). An introduction to Bernoullian Utility Theory: I utility functions. *The Swedish Journal of Economics*, *69*, 163–183.
- Jessup R. K., Bishara A. J., Busemeyer J. R. (2008). Feedback produces divergence from prospect theory in descriptive choice. *Psychological Science*, *19*, 1015–1022.
- Johnson, R. A., & Wichern, D. W. (1998). *Applied multivariate statistical analysis*. Upper Saddle River, NJ: Prentice Hall.
- Kahneman, D. & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review*, *103*, 582–591.
- Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioural economics. *The American Economic Review*, *93*, 1449–1475.
- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 49–81). New York: Cambridge University Press.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–292.
- Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review*, *103*, 582–591.
- Kahneman, D., Slovic, P., & Tversky A. (1982). (Eds.), *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.
- Karni, E., Schmeidler, D., & Vind, K. (1983). On state dependent preferences and subjective probabilities. *Econometrica*, *51*, 1021–1031.
- Kass, R. E., & Raftery, A. E. (1995). Bayes Factors. *Journal of the American Statistical Association*, *90*, 773–795.
- Kowler, E., & Blaser, E. (1995). The accuracy and precision of saccades to small and large targets. *Vision Research*, *35*, 1741–1754.
- Lee, M. D., & Wagenmakers, E.-J. (2005). Bayesian statistical inference in psychology: Comment on Trafimow (2003). *Psychological Review*, *112*, 662–668.

- Lyons, J., Hansen, S., Hurding, S., & Elliott, D. (2006). Optimizing rapid aiming behaviour: movement kinematics depend on the cost of corrective modifications, *Experimental Brain Research*, *174*, 95–100.
- Maddox, T. W. (2002). Toward a unified theory of decision criterion learning in perceptual categorization. *Journal of the Experimental Analysis of Behavior*, *78*, 567–595.
- Maloney, L. T., & Mamassian, P. (2009). Bayesian decision theory as a model of visual perception: Testing Bayesian transfer, *Visual Neuroscience*, *26*, 147-155.
- Maloney, L. T., Trommershäuser, J., & Landy, M. S. (2007). Questions without words: A comparison between decision making under risk and movement planning under risk. In W. Gray (Ed), *Integrated Models of Cognitive Systems* (297–313). New York: Oxford University Press.
- Martinez, W. L., & Martinez, A. R. (2002). *Computational statistics handbook with Matlab*. London: Chapman & Hall.
- Mates, J., Radil, T, Mueller, U., & Poeppel E. (1994). Temporal integration in sensorimotor synchronization. *Journal of Cognitive Neuroscience*, *6*, 332 – 340.
- Mosteller, F., & Nogee, F. (1951). An Experimental Measurement of Utility. *Journal of Political Economy*, *59*, 371–404.
- Nakagaki, T., Yamada, H., Toth, A. (2000). Maze-solving by an amoeboid organism, *Nature*, *407*, 470.
- Navalpakkam V., Koch C., Rangel A., & Perona, P. (2010). Optimal reward harvesting in complex perceptual environments, *Proceedings of the National Academy of Sciences: USA*, *107*, 5232–5237.
- Navalpakkam, V., Koch, C., & Perona, P. (2009). Homo economicus in visual search. *Journal of Vision*, *9*, 1–16.
- Neumann, von J., & Morgenstern, O. (1953). *Theory of games and economic behavior*. 3rd Ed. Princeton: Princeton University Press.
- Nickerson, R. S. (2008). *Aspects of rationality: Reflections on what it means to be rational and whether we are*. Hove: Psychology Press.
- Ozer, D. J. (1987). Personality, intelligence and spatial visualization: Correlates of mental rotations test performance. *Journal of Personality and Social Psychology*, *53*, 129–134.
- Parmigiani, G., & Inoue, L. Y. T. (2009). *Decision theory: Principles and approaches*. Chichester, UK: John Wiley & Sons.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics:

- Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Ploner, C. J., Ostendorf, F., & Dick, S. (2004). Target size modulates saccadic eye movements in humans. *Behavioral Neuroscience*, *118*, 237–242.
- Pompilio, L., & Kacelnik, A. (2010). Context-dependent utility overrides absolute memory as a determinant of choice. *Proceedings of the National Academy of Sciences USA*, *107*, 508-512.
- Prelec, D. (1998). The probability weighting function. *Econometrica*, *66*, 497–527.
- Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, *68*, 1281-1292.
- Rabin, M. , & Thaler, R. H. (2001). Anomalies: Risk Aversion. *Journal of Economic Perspectives*, *15*, 219-232.
- Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, *23*, 1-14.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*, 225–237.
- Savage, L. J. (1954/1972). *The foundations of statistics (2nd Ed)*, New York: Dover Publications.
- Schaik, van P., Kusev, P. & Juliusson, A. (2011). Human preferences and risky choices. *Frontiers in Psychology*, *2*, 1-2.
- Schmidt, R. A., Zelaznik, H. N., Hawkins, B., Frank, R.S., & Quinn, J. T. (1979). Motor-output variability: A theory for the accuracy of rapid motor acts. *Psychological Review*, *86*, 415–451.
- Schouten, J. F., & Bekker, J. A. M. (1967). Reaction time and accuracy, *Acta Psychologica*, *27*, 143–153.
- Schrater, P. R., & Kersten, D. (2002). Vision, psychophysics, and Bayes. In: (R. P. N. Rao, B. A. Olshausen, & M. S. Lewicki, Eds.). *Probabilistic models of the brain: Perception and neural function*, MIT Press.
- Searle, J. R., (2001). *Rationality in action*. Cambridge, Mass: MIT Press.
- Seydell, A., McCann, B. C., Trommershäuser, J., & Knill, D. C. (2008). Learning stochastic reward distributions in a speeded pointing task. *The Journal of Neuroscience*, *28*, 4356–4367.
- Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in

- the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86, 1916–1936.
- Shafir, S., Reich, T., Tsur, E., Erev, I., & Lotem, A. (2008). Perceptual accuracy and conflicting effects of certainty on risk-taking behaviour. *Nature*, 453, 917–921.
- Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, 15, 233 – 250.
- Shepard, R. N., & Metzler J. N. (1967). Mental rotation of three-dimensional objects. *Science*, 171, 701–703.
- Shimozaki, S., Kingstone, A., Olk, B., Stowe, R., & Eckstein, M. (2006). Classification images of two right hemisphere patients: A window into the attentional mechanisms of spatial neglect. *Brain Research*, 1080, 26–52.
- Simon, H. A. (1959). Theories of decision-making in economic and behavioral science. *The American Economic Review*, 49, 253–283.
- Smith, V. L., (1976). Experimental economics: Induced value theory. *The American Economic Review*, 66, 274 - 279.
- Stewart, N., Chater, N., Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, 53, 1–26.
- Stott, H. P., (2006). Cumulative prospect theory's functional menagerie. *Journal of Risk and Uncertainty*, 32, 101-130.
- Sutherland, S. (2007). *Irrationality*. London: Pinter & Martin Ltd.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT Press: Cambridge, MA.
- Tanaka, H., Krakauer, J. W., & Qian, N. (2006). An optimization principle for determining movement duration. *Journal of Neurophysiology*, 95, 3875–3886.
- Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics*, 112, 647–661.
- Thorndike, E. L. (1911). *Animal Intelligence*. New York: Macmillan
- Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5, 1226–1235.
- Trommershäuser, J., Gepshtein, S., Maloney, L. T., Landy, M. S., & Banks, M. S. (2005). Optimal compensation for changes in task-relevant movement variability. *The Journal of Neuroscience*, 25, 7169–7178.
- Trommershäuser, J., Landy, M. S., & Maloney, L. T. (2006). Humans rapidly

- estimate expected gain in movement planning. *Psychological Science*, *17*, 981–988.
- Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003a). Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America A*, *20*, 1419–1432.
- Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003b). Statistical decision theory and trade-offs in the control of motor response. *Spatial Vision*, *16*, 255–275.
- Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences*, *12*, 291–297.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*, 453–458.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.
- Ulehla, Z. J. (1966). Optimality of perceptual decision criteria. *Journal of Experimental Psychology*, *71*, 564–569.
- Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighed when rare outcomes are experienced (rarely)? *Psychological Science*, *20*, 473–479.
- van Beers R. J., Haggard P., Wolpert D. M. (2004). The role of execution noise in movement variability. *Journal of Neurophysiology*, *91*, 1050–1063.
- von Neumann, J. , & Morgenstern, O. (1944/1955). *Theory of games and economic behavior*. Princeton: Princeton University Press
- Weibull, W. (1951). A statistical distribution function of wide applicability. *Journal of Applied Mechanics*, *18*, 293–297.
- Whiteley, L., & Sahani, M. (2008). Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes. *Journal of Vision*, *8*, 1–15.
- Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & Psychophysics*, *63*, 1293–1313.
- Winterfeldt, D, von, & Edwards, W. (1982). Costs and payoffs in perceptual research. *Psychological Bulletin*, *91*, 609–622.
- Wu, S-W., Delgado, M. R., & Maloney, L. T. (2009). Economic decision-making compared with an equivalent motor task. *Proceedings of the National Academy of Sciences USA*, *106*, 6088–6093.

- Wu, S-W., Delgado, M. R., & Maloney, L. T. (2011). The neural correlates of subjective utility of monetary outcome and probability weight in economic and in motor decision under risk. *The Journal of Neuroscience*, *31*, 8822-8831.
- Wu, S-W., Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2006). Limits to human movement planning in tasks with asymmetric gain landscapes. *Journal of Vision*, *6*, 53–63.