Personality and Individual Differences 51 (2011) 242–247

Contents lists available at ScienceDirect

Personality and Individual Differences

journal homepage: www.elsevier.com/locate/paid

An economic perspective on the Reinforcement Sensitivity Theory of personality

Phillip J. Hall [a], Chew Wuei Chong [a], Neil McNaughton [a,*], Philip J. Corr [b]

[a] Department of Psychology and Neuroscience Research Centre, University of Otago, Dunedin, New Zealand
[b] School of Social Work and Psychology and the Centre for Behavioural and Experimental Social Science (CBESS), University of East Anglia, UK

ARTICLE INFO

Article history:
Available online 14 July 2010

Keywords:
Reinforcement Sensitivity Theory
Neuroeconomics
Reward
Punishment
Gain
Loss
Approach
Avoidance

ABSTRACT

Reinforcement Sensitivity Theory postulates personality factors of 'reward sensitivity' and 'punishment sensitivity' linked to neural systems that control approach and avoidance, respectively. In contrast, behavioural economics distinguishes gain ('reward') and loss ('punishment') valuation systems that are orthogonal to approach/avoidance behaviour. We combined gain and loss with both their presentation and omission and found evidence for separate gain valuation, loss valuation, approach, and avoidance systems. This suggests that it is possible to integrate valuation/input and behaviour/output views of 'reward' and 'punishment' in a way that may be of use to both personality theory and economics and so forge closer links between these two major perspectives on decision-making and behaviour.

1. Introduction

The aim of this article is to provide a basis for integrating neurally-based personality theories, such as Reinforcement Sensitivity Theory (RST), with neuroeconomic findings and theories by resolving an apparent ambiguity in the terms 'reward' and 'punishment'.

The RST of personality (for overview see Corr, 2008) sees factors of 'reward sensitivity' and 'punishment sensitivity' as mediated by neural systems that control approach and avoidance. For RST, omission/termination of expected reward is a punishment – leading to avoidance/escape – and vice versa for punishment (Gray, 1975, 1982; Gray & McNaughton, 2000).

Neuroeconomics has demonstrated distinct gain and loss valuation systems (Seymour, Daw, Dayan, Singer, & Dolan, 2007; Yacubian et al., 2006), and behavioural economics has shown that the loss system is stronger (Tversky & Kahneman, 1991). That is, gain and loss omission can both produce approach; but, for the same dollar value, loss produces the stronger effects. Since the approach response is the same in the two cases, loss aversion cannot be due to a difference between the approach and avoidance systems underlying RST.

Fig. 1 shows the interaction of valuation input and behaviour output systems when loss and gain are the consequences of responding or not responding. This separation of input and output components of the system, and the subtractive relationship between approach and avoidance output tendencies, is the same as in the most detailed previous formulation of state aspects of RST (Gray & Smith, 1969). What is different is that valuation precedes the effect of omission in the computational chain. Without this, loss aversion would not occur in the case of omission. Loss from an existing store (of food pellets as opposed to dollars) has not often been manipulated in the animal behavioural experiments on which RST is based. This may account for its neglect in previous versions of the theory.

In Fig. 1, the use of dollars allows external real values of positive and negative consequences to be directly compared as they are on the same scale. In the figure all consequences for an individual trial are set to a value of $1. The economic literature tells us that these identical external $1 values will not all generate the same internal motivational value. On average, people respond as though the internal value of a loss is greater (110% in the figure) than that of a gain (90%).

Knowing the value of the expected loss or gain does not tell us whether approach or avoidance will be produced. The probability or rate of a response will be increased not only if it produces gain but also if it prevents loss that would otherwise have occurred. Conversely, responding will be decreased when it results in loss or gain omission. This is represented in Fig. 1 by the crossed connections between gain and loss valuation and approach and avoidance tendencies for the cases where responses result in omission of the consequence of inaction.

The important point to note here is that the distinction between gain and loss valuation is orthogonal to the distinction between approach and avoidance – and the objects of each can in both cases be referred to as 'reward' and 'punishment' depending on your precise use of these terms. The loss aversion reported in economic

* Corresponding author. Address: Department of Psychology, University of Otago, P.O. Box 56, Dunedin, New Zealand. Tel.: +64 3 479 7643; fax: +64 3 479 8335.
E-mail address: nmcn@psy.otago.ac.nz (N. McNaughton).

0191-8869/$ - see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.paid.2010.06.023

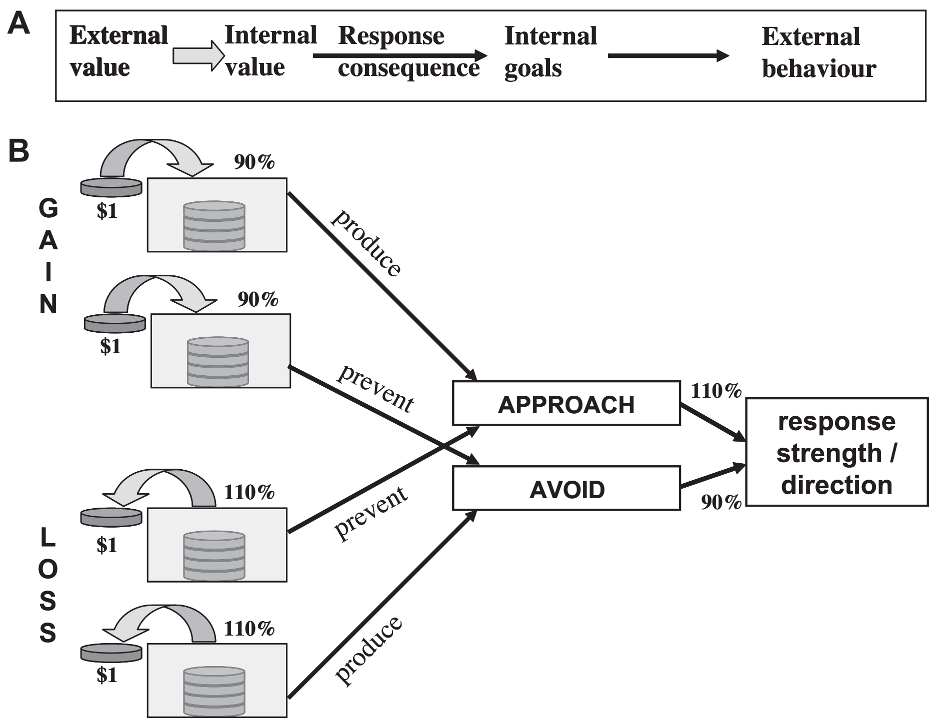

Fig. 1. The combination of valuation and operant factors that determines response strength and direction. Items with a specific external value ($1) that can be gained or lost are represented by a particular internal value that will depend both on the level of "hunger" for the item and on whether it is gained or lost. Economic analysis has shown that the same external value generally has a greater internal value if it is a loss (110%) than if it is a gain (90%). The effect of this internal valuation on behaviour depends on the consequences of responding. Gain production and loss prevention activate approach; loss production and gain prevention activate avoidance. Concurrent approach and avoidance tendencies are then integrated to determine the direction and strength of responding. Approach and avoidance will have different strengths (e.g., 110% and 90%) depending both on factors of reinforcement sensitivity and on the distance from goal achieved by responding. (Approach and avoidance have different goal gradients, see text.)

experiments is, potentially, the result of the operation of quite different personality factors than those that RST sees as specifically controlling approach and avoidance. As can be seen from Fig. 1, taking account of gain and loss valuation sensitivities does not eliminate the need to also take into account approach and avoidance sensitivities. Approach and avoidance tendencies (labelled internal goals in the figure) can and usually are concurrently activated to some extent and it is the result of their interaction that we observe as variations in the intensity and direction of behaviour.

Previous experiments have not assessed both gain/loss sensitivity differences and distinct approach/avoidance sensitivity differences within the same basic paradigm. We, therefore, administered a task in which human participants received all four of the basic conditions illustrated in Fig. 1.

In one phase of the experiment, participants started with zero dollars and gain was pitted against loss: clicking on a target resulted in a 50:50 chance of gaining money or losing money. Nonresponding resulted in no gain or loss. People responded more on trials with net gain than those with net loss and so the final amount earned was always positive. The gain and loss values were fixed and known within a block of trials. Across blocks, gain and loss values were changed. Gain and loss each had the same set of possible values and all possible combinations of these gain and loss values were tested in combination. The response to be made was modelled on rat runway experiments and involved moving the mouse cursor a fixed distance to a target and clicking on the target. The overall speed of responding was measured. Speed was used as a measure rather than simple choice as it gives a numerical measure of motivation on each trial and is the measure usually used in the animal experiments on which RST is based.

In a second phase, all details were the same except that participants started with a set number of dollars and each click prevented the loss of money but with a 50:50 risk of preventing the gain of money. That is, they received the "prevent" consequences of Fig. 1. (Note that from a "rational economic" perspective there is no difference between this second set of phases and the first. A gain of $1 and prevention of loss of $1 have identical consequences on take-home amount. The click response differs only in whether it is against a background of starting the experiment with nothing or with something.)

Across these conditions, we then analysed sensitivity to dollar consequences either: (1) with respect to approach and avoidance (averaging across gain and loss); or (2) with respect to the difference between gain and loss (averaging across approach and avoidance).

2. Method

2.1. Participants

Twenty-one (9 male and 12 female) University of Otago students were recruited through Student Job Search. They ranged in age between 18 and 45 years old (mean age = 24.5). They were informed that they would receive a guaranteed minimum amount of NZ$9 (the legal minimum wage at the time) for their participation, and could earn up to NZ$15 based on their performance. The task difficulty was sufficiently low that all participants in this experiment received the full payment. It should be noted that since speed was our measured variable (and not accuracy) this payment result does not represent a ceiling effect.

2.2. Design

A within-participants, repeated-measures, design was used with Reaction Time (RT) as the dependent variable. There were four experimental phases, delivered in the following order: (1)

pure gain; (2) pure loss omission; (3) mixed gain/loss; and, (4) mixed loss omission/gain omission. Each of the phases was experienced twice: once with descending experimental dollar (E$) values and once with ascending E$ values. The pure gain phase and the pure loss omission phase each had nine levels varying from E$0 to E$8 in E$1 steps. Phase 3 had the full factorial combination of five levels (E$0–E$8 in steps of E$2) of gain combined with five levels (E$0–E$8 in steps of E$2) of loss. Phase 4 was like phase 3, except that responding prevented loss or prevented gain omission). Thus, the two possible outcomes in phase 3 are similar (except in the range of monetary values) to the top and bottom ones depicted in Fig. 1, while the outcomes in phase 4 are similar to the two middle ones.

2.3. Procedure

Participants were told that by choosing whether or not to click a computer mouse they would earn or lose the experimental dollars shown on a computer screen within a blue box on each trial. The amounts earned could be exchanged at the end of the experiment, for New Zealand dollars at the exchange rate of 100 experimental dollars per New Zealand dollar, to a maximum of NZ$15. Participants were given a block of practice trials prior to each of the four phases consisting of a number of trial blocks (18 in phases 1 and 2, 50 in phases 3 and 4), with each trial block five trials long. An information screen at the start of each trial block informed participants of the monetary consequences for the upcoming block; and signalled whether responding would produce or prevent the consequences.

In all phases of the task, a blue box appeared at an unpredictable location on the screen and the participant needed to move the mouse icon on the screen, using the computer mouse, to the box and click on it to obtain the programmed consequence. Both of the two possible consequences for that trial were displayed in the box. If no response was made within 3000 ms, the trial terminated without consequences. While the location of the box was controlled by a random number generator, it always occurred at a fixed distance from the previous box, and the mouse pointer was re-set back to the location of the previous box, so that the distance that the mouse pointer would have to move would be the same in every trial. The mouse pointer was inactivated between trials. In all phases, the actual consequence of pressing, or not pressing, was signalled both by text indicating the actual amount of change (including E$0 change) in experimental dollars received and by a change in the colour of the blue box: to green for a gain in experimental dollars, red for a loss, and grey for no change.

In the first phase, participants started with no experimental dollars, and could earn money by clicking on the blue box. The amount that could be gained varied sequentially from E$8 to E$0, then E$0–E$8 in E$1 steps. The box changed colour to green when they clicked the box to indicate they had earned the amount of experimental dollars shown within the blue box, or to grey to indicate they had not if they failed to respond within the 3 s response period.

In the second phase, participants started with E$360 and could prevent loss of money by clicking on the blue box. The box colour changed to grey to indicate they had prevented the loss of experimental dollars shown within the blue box or to red to indicate they had received the loss, if they did not respond within the 3 s response period. The amount that could be lost varied sequentially across trial blocks from E$8 to E$0, then E$0–E$8 in E$1 steps.

In phase 3, participants started with E$180. By choosing to click, there was a 50% chance that they would gain the positive amount of experimental dollars, and a 50% chance that they would lose the negative amount of experimental dollars, both shown within the blue box. If participants chose not to click, then they would not gain or lose any experimental dollars. The amount that could be gained or lost was varied from E$8 to E$0 in $2 steps, and all combinations of gain and loss value were tested twice: once each in descending and then ascending order.

In phase 4, participants started with E$180. This phase was identical to phase 3, except that this time by choosing to click, there was a 50% chance that they would prevent the loss of experimental dollars, and a 50% chance that they would prevent the gain of experimental dollars. That is, if participants chose to click, then they would not gain or lose any experimental dollars; and if they did not click they would obtain the programmed consequences.

In both phases 3 and 4, the blue box would change colour to green if they gained experimental dollars, to red if they lost experimental dollars and to grey if there was no change in net experimental dollars.

2.4. Analysis

Data for phase 1 and phase 2 were submitted to one nested repeated-measures analysis of variance; and those of phase 3 and phase 4 were submitted to a second nested repeated-measures analysis of variance. All the factors detailed under Section 2.2, above, were extracted with orthogonal polynomial components fitted to dollar factors. The analyses were carried out using the Genstat statistical package.

The data from phase 1 and phase 2 did not show a simple proportionality of speed of response to amount of reinforcement. Likewise, the data from phase 3 and phase 4 did not show simple additivity of gains and losses. Instead, the data were as would be expected from the matching law, which holds for magnitude of reinforcement as well as for the more usual variation in rate of delivery of a fixed size of reinforcement (De Villiers, 1977).

A simple form of the matching law that has been assessed with the combination of positive and negative reinforcement (Farley, 1980) is:

$$\text{Response rate}_1 = k(a_1 - cr_1)/((a_1 - cr_1) + (a_2 - cr_2) + r_o) \quad (1)$$

where a is the total attraction value ('reward', gain or loss omission), r is the total repulsion value ('punishment', loss or gain omission), c a constant reflecting the "exchange" rate between the units of measurement of attraction and repulsion, r_o is the unprogrammed reinforcement of other behaviours, k is the nominal maximum response rate and the subscripts refer to two distinct concurrently available manipulanda.

For the current experiments, we were measuring speed of an extended response rather than rate of a discrete response and we were directly reinforcing and measuring only a single response, with all competition coming from other behaviours. This reduces Eq. (1) to:

$$\text{speed} = k(a - cr)/((a - cr) + r_o) \quad (2)$$

We made two further modifications to Eq. (2) to fit the circumstances of the current experiment. First, since the external units of attraction and repulsion were the same (i.e., dollars), and since RST deals with separate sensitivities of the approach and avoidance systems, we replaced the single exchange rate constant, c, with two "exchange rate" constants: one (A) for the rate between dollars and the effect of attraction on speed; the other (R) for the rate between dollars and the effect of repulsion on speed. Second, there was clear responding at zero dollar value indicating that there was some small intrinsic gain from responding that overcame the small intrinsic loss that must be assumed to be incurred by responding. Therefore, to the existing (experimenter determined) extrinsic gain and loss values, we added intrinsic gain and loss parameters. The resultant equation is:

$$\text{speed} = k(A(a) - R(r))/(A(a) + R(r) + r_o) \quad (3)$$

where $a = a_{\text{intrinsic}} + a_{\text{extrinsic}}$ is the total attraction value (gain or loss omission), $r = r_{\text{intrinsic}} + r_{\text{extrinsic}}$ is the total repulsion value (loss or gain omission), r_o is the unprogrammed reinforcement of other behaviours, k is the nominal maximum speed, A is the attraction sensitivity and R is the repulsion sensitivity.

Extrinsic values are those imposed by the experimenter and are known. The two intrinsic values ($a_{\text{intrinsic}}$, $r_{\text{intrinsic}}$) and the parameters k, A, R and r_o must be calculated by a least squares fit to the data. This was carried out using the solver system in Microsoft® Office Excel 2003 individually for each participant's data and also for the overall average across participants. While there are a variety of complexities that can be introduced into such equations (Killeen, 1994), the basic form of Eq. (2) has the advantage of simplicity and the parameters we have added in Eq. (3) are required by the observed behaviour. Importantly (see Section 3) good, and significant, fits were obtained. Adding further parameters would not have greatly improved the fits in most cases and would have potentially generated a problem for the testing of significance as the number of parameters would have been large with respect to the number of data points being fitted.

3. Results

3.1. Phase 1 and phase 2

In neither phase 1 nor phase 2 was there a simple monotonic relationship between dollar value and responding. There was moderate responding at zero value and essentially asymptotic responding at all other values (data not shown).

3.2. Phase 3 and phase 4 – overall ANOVA

Initial analysis of phase 3 and phase 4 can be seen as testing concurrently for effects that differ between approach and avoidance (averaged across gain and loss valuations) and for differences between gain and loss (averaging across approach and avoidance) and, via the interaction terms, for independence of these two types of effect. These effects were significant (direction × gain value × loss value: dev.quad.lin, $F(1, 1520) = 5.83$ $p = 0.016$; dev.lin.quad, $F(1, 1520) = 11.77$, $p < 0.001$; dev.cub.quad, $F(1, 1520) = 9.73$, $p = 0.002$).

3.3. Phase 3 and phase 4 – approach versus avoidance

Fig. 2A shows the variation in response speed with changes in attraction (gain/loss omission) and repulsion (loss/gain omission) value, i.e., averaged across gain and loss, eliminating any effects of loss aversion.

Important points to note about these results are that: (1) a net zero dollar value for the making of a response does not result in zero response (i.e., minimal speed); so, as in phase 1 and phase 2, there appears to be an intrinsic value to responding even for no dollars and this outweighs the small intrinsic response cost that must also exist; (2) a net $2 difference between attractor and repulsor values produces speeds that depend on their absolute values rather than having a fixed effect depending on the net value; (3) the interaction of attractor and repulsor values produces a curvilinear relationship (see especially the curve for $2 attractor value).

The result of fitting the value function we derived from the matching law (see Section 2) to all the points concurrently is shown in Fig. 2B. This fit accounted for 98% of the variance of this averaged data set. Fitting the same function to individual participant data accounted for 89% of the variance on average (range 49–95%). With six free parameters, fitting 25 observed values in each case, the critical value at $p < 0.05$ is 75% with only one participant's data (49%) failing this criterion; while at $p < 0.01$, it is 81% with 19/20 participants' data passing this criterion. The goodness of these fits shows reasonable generalisation of the previous animal analyses to this human task.

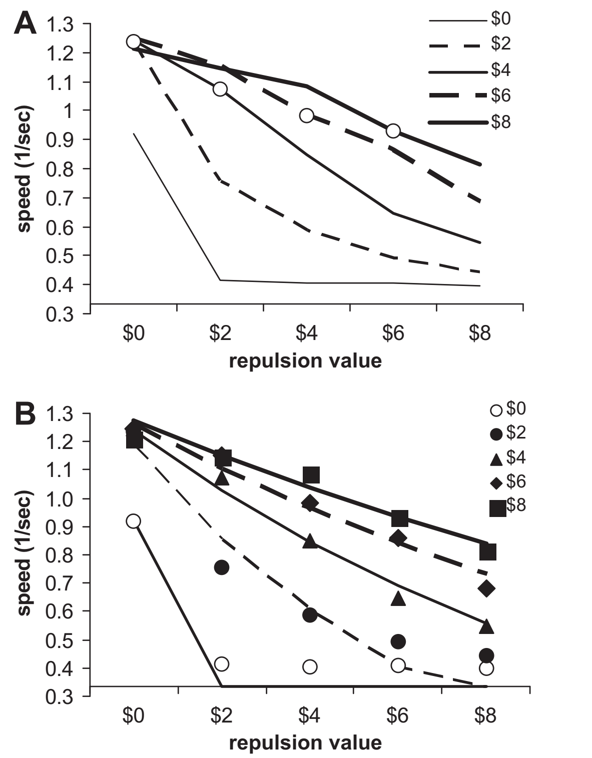

Fig. 2. (A) Observed speeds resulting from the combination of a specific attraction value for the response with a particular repulsion value. Attraction dollar value is gain averaged with omission of loss (separate values plotted as separate curves). Repulsion dollar value is loss averaged with the omission of gain and is plotted on the *x*-axis. The probability of gain or loss on any particular trial was equal. Open circles indicate the point on each curve at which the net attraction value averages to $2. (B) The same data represented as point values with the curves resulting from a single optimised fitted function based on previous animal behaviour analysis (see text). Note the flat-line values at the bottom of the graph and the placing of the *x*-axis intercept on the *y*-axis are the result of the restriction (both in reality and on the fitting function) of a fixed 3 s maximum response time.

Perhaps the most important aspect of these results is that, for all participants tested, individual approach strength was always greater than individual avoidance strength (ranging from 1.6 to 5.4 times). (We excluded the participant with the 49% variance data fit. Their ratio, 10.43, is likely to be an artefact of the poor fit.) The approach and avoidance systems clearly have different dollar sensitivities. However, as detailed in the discussion, the values obtained here need to be treated with caution as approach and avoidance tendencies have different goal gradients.

3.4. Phase 3 versus phase 4 – gain versus loss

If we see approach as linked to gain, and avoidance as linked to loss, then our finding of a stronger approach tendency than avoidance tendency seems to contradict the neuroeconomic literature on loss aversion. However, when we look at the *difference* between the use of gain and the use of omission of loss to promote responding and the *difference* between the omission of gain and presentation of loss to inhibit responding (i.e., the differences between the pairs of curves that were averaged to produce Fig. 2) then we are looking at differences in gain and loss *valuation* independent of approach and avoidance responses.

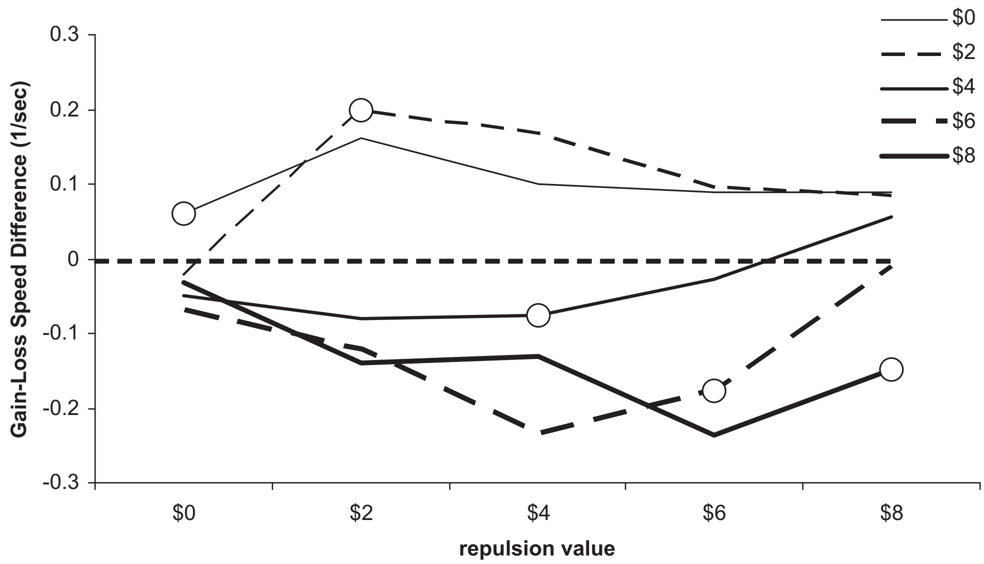

Fig. 3. Observed *differences* in speed between gain and loss manipulations for the combination of a specific attraction value for the response with a particular repulsion value. Attraction dollar value is gain averaged with omission of loss (separate values plotted as separate curves). Repulsion dollar value is loss averaged with the omission of gain and is plotted on the *x*-axis. The probability of gain or loss on any particular trial was equal. Open circles indicate the point on each curve at which the attraction and repulsion values are equal.

If we assess gain and loss valuation, independent of approach and avoidance, we find (Fig. 3) that we have evidence for loss aversion, e.g., speed is generally greater when responding is driven by omission of loss rather than presentation of gain. This effect occurs mainly with higher values of loss and is easiest to see (open circles in the figure) where attraction and repulsion values are equal.

4. Discussion

The main findings of the present experiment are: (1) approach and avoidance tendencies have different strengths for a range of given external dollar values when averaged across gain and loss omission or loss and gain omission, respectively; and, (2) at least with higher values, the loss aversion typically reported in the literature is obtained when averaging across approach and avoidance tendencies.

The distinction between gain/loss evaluation and approach/avoidance behaviour has been conflated within the previous use of the terms 'reward' and 'punishment', especially in the RST literature. Integration of the distinct behavioural and valuation perspectives of reward and punishment should be to the mutual benefit of personality psychology and current developments in behavioural economics that attempt to provide objective neural anchors for economic rules (Glimcher, Dorris, & Bayer, 2005; Glimcher & Rustichini, 2004; Sanfey, 2007; Sanfey, Loewenstein, McClure, & Cohen, 2006; van't Wout, Kahn, Sanfey, & Aleman, 2006; Zak, 2004).

The key point is that valuation of gain and loss (which economics tells us will involve loss aversion) must precede their effect on approach and avoidance (which previous animal research tells us involve different neural systems with different sensitivities). Whether valuation is followed by approach or avoidance then depends on consequences: presentation or omission (Fig. 1).

Strictly, RST should see its critical personality factors as being related to the variations illustrated in Fig. 2 and not Fig. 3. RST posits a behavioural approach system pitted against systems that result in avoidance or suppression of approach. One implication of the current results is that experiments assessing the approach and avoidance systems should ensure that gain and loss are balanced by loss omission and gain omission if loss aversion is not to confound estimates of approach and avoidance tendencies. Likewise, variations in loss aversion can be assessed, averaging across approach and avoidance as in Fig. 3 and will potentially demonstrate the operation of gain and loss sensitivity as personality factors. Whether repulsion or loss maps better to any particular current personality measure (such as BIS sensitivity, or Harm Avoidance) requires further research.

It might seem that our results have demonstrated a contrast between approach preference and loss aversion. However, it is important to note here that, in typical animal experiments, approach and avoidance gradients are different (Gray, 1975, pp. 234–236) and so the measured relative strength of approach and avoidance will change depending on distance from a goal as well as with more stable personality factors. Approach will be stronger further from the goal, and a sufficiently strong avoidance tendency will be stronger than the approach tendency closer to the goal – with a cross-over at intermediate distances. In the present case, then, it is possible that the ultimate goal (receipt of money at the end of the experiment) is relatively distant–giving rise to a greater relative strength of approach, which would not be obtained with more immediate and direct gain or loss.

This does not invalidate our distinction between approach/avoidance and gain/loss. Our calculations deriving these values were carried out on the same data set and so necessarily at the same distance from the ultimate goals. We can, therefore, take the fact that approach was greater than avoidance as being inconsistent with the control of avoidance being identical to the control of loss (given our demonstration of loss aversion); and, of course, any effect of loss *per se* was averaged out before fitting the approach–avoidance curves.

Perhaps the most important point for practical purposes in the assessment of personality sensitivities (whether approach/avoidance, or gain/loss) is that, at least with the current measures, effects of variation in the strength of approach tendencies can only be assessed in the presence of some avoidance tendency. As found in our phases 1 and 2, if there is insufficient avoidance then speed goes rapidly to asymptote and demonstrates no effects of variation in the strength of the approach tendency. The sensitivity of approach and avoidance tendencies cannot, therefore, be measured directly.

While absolute approach sensitivity and absolute avoidance sensitivity cannot be reliably estimated, the curves that result when they are systematically varied can be fitted, with a high degree of accuracy using the matching law. These equations allow the *ratio* of approach and avoidance strength to be estimated and so their relative strengths to be determined. We do not claim that the specific mathematical form of the equation we used is correct (Killeen, 1994) – but the goodness of fit gives us reason to suppose that alternatives to it would give the same fundamental conclusions about relative strengths at the particular point along the approach and avoidance gradients obtained in this experiment.

For personality assessment, the approach:avoidance strength ratio could probably be fractionated, via multiple regression, to assess unique components of variance in the ratio determined by existing personality scales – but the interpretation of the results will be non-trivial given the capacity of the approach and avoidance systems to interact (Gray & Smith, 1969). To obtain real anchors for the independent sensitivities, one would probably need to go to fMRI or related cellular techniques (Sanfey, 2007).

It should be noted here that our problem in resolving absolute approach and avoidance sensitivities recurs with loss aversion. That is, loss aversion is measured relative to gain and is, in that sense, simply a ratio or difference measure. Resolving factors related to gain sensitivity and loss sensitivity in absolute terms, then, represents the same problems as for approach and avoidance – with fMRI or EEG techniques providing the most obvious way of solving the problem.

Our paradigm and its results are only a very preliminary attempt to forge closer theoretical and empirical bonds between animal-based theories of human personality (e.g., RST) and principles and findings from behavioural and neural economics. However, we believe that this admittedly preliminary attempt has already identified a potentially fundamental distinction in human preferences under conditions of 'reward' and 'punishment': the distinction between gain/loss evaluation and approach/avoidance behaviour. This 2×2 classification deserves further scrutiny, for it may hold significant implications for how individual differences in the sensitivity to 'reward' and 'punishment' are conceptualised and measured, as well as providing a framework for the integration of behaviourist with economic data. Our results also appear to show that direct behavioural measurement of specific single sensitivities will be difficult or impossible. This suggests that, as with the current developments in neuroeconomics, personality theory will benefit from the use of neural measures. Critically, neural measures provide a clear common ground for linking biological theories of personality with neuroeconomics.

References

Corr, P. J. (Ed.). (2008). *The Reinforcement Sensitivity Theory of Personality*. Cambridge: Cambridge University Press.

De Villiers, P. (1977). Choice in concurrent schedules and a quantitative formulation of the law of effect. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behaviour* (pp. 233–278). Englewood Cliffs: Prentice Hall.

Farley, J. (1980). Reinforcement and punishment effects in concurrent schedules: A test of two models. *Journal of the Experimental Analysis of Behavior, 33*, 311–326.

Glimcher, P. W., Dorris, M. C., & Bayer, H. M. (2005). Physiological utility theory and the neuroeconomics of choice. *Games and Economic Behaviour, 52*, 213–256.

Glimcher, P. W., & Rustichini, A. (2004). Neuroeconomics: The consilience of brain and decision. *Science, 306*, 447–452.

Gray, J. A. (1975). *Elements of a two-process theory of learning*. London: Academic Press.

Gray, J. A. (1982). *The Neuropsychology of Anxiety: An enquiry in to the functions of the septo-hippocampal system*. Oxford: Oxford University Press.

Gray, J. A., & McNaughton, N. (2000). *The Neuropsychology of Anxiety: An enquiry into the functions of the septo-hippocampal system*. Oxford: Oxford University Press.

Gray, J. A., & Smith, P. T. (1969). An arousal-decision model for partial reinforcement and discrimination learning. In R. Gilbert & N. S. Sutherland (Eds.), *Animal discrimination learning* (pp. 243–272). London: Academic Press.

Killeen, P. R. (1994). Mathematical principles of reinforcement. *Behavioral and Brain Sciences, 17*, 105–172.

Sanfey, A. G. (2007). Decision neuroscience: New directions in studies of judgment and decision making. *Current Directions in Psychological Science, 16*, 151–155.

Sanfey, A. G., Loewenstein, G., McClure, S. M., & Cohen, J. D. (2006). Neuroeconomics: Cross-currents in research on decision-making. *Trends in Cognitive Sciences, 10*, 108–116.

Seymour, B., Daw, N., Dayan, P., Singer, T., & Dolan, R. (2007). Differential encoding of losses and gains in the human striatum. *Journal of Neuroscience, 27*, 4826–4831.

Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference dependent model. *The Quarterly Journal of Economics, 106*, 1039–1061.

van't Wout, M., Kahn, R. S., Sanfey, A. G., & Aleman, A. (2006). Affective state and decision-making in the Ultimatum Game. *Experimental Brain Research, 169*, 564–568.

Yacubian, J. L., Glascher, J. N., Schroeder, K. T., Sommer, T. B., Braus, D. F. E. F., & Buchel, C. R. (2006). Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *Journal of Neuroscience, 26*, 9530–9537.

Zak, P. J. (2004). Neuroeconomics. *Philosophical Transactions of the Royal Society of London B, 359*, 1737–1748.