Ageing disrupts reinforcement learning whilst learning to help others is preserved

Jo Cutler*[a,b,c], Marco Wittmann[a,b,], Ayat Abdurahman[a,b,d], Luca Hargitai[a], Daniel Drew[a,b], Masud Husain[a,b], Patricia Lockwood*[a,b,c]

[a]Department of Experimental Psychology, University of Oxford, United Kingdom
[b]Wellcome Centre for Integrative Neuroimaging, Department of Experimental Psychology, University of Oxford, United Kingdom
[c]Centre for Human Brain Health, School of Psychology University of Birmingham, United Kingdom
[d]Department of Psychology, University of Cambridge, United Kingdom

*Correspondence should be addressed to:
Jo Cutler: J.L.Cutler@bham.ac.uk
Centre for Human Brain Health, Birmingham University, Edgbaston Park Road, Birmingham, B15 2TT, United Kingdom.

Patricia Lockwood: patricia.lockwood@psy.ox.ac.uk
Department of Experimental Psychology, University of Oxford, Tinsley Building, Mansfield Road, Oxford, United Kingdom OX1 3SR.

Abstract

Reinforcement learning is a fundamental mechanism displayed by many species. However, adaptive behaviour depends not only on learning about actions and outcomes that affect ourselves, but also those that affect others. Here, using computational reinforcement learning models, we tested whether young (age 18-36) and older (age 60-80, total n=152) adults can learn to gain rewards for themselves, another person (prosocial), or neither individual (control). Detailed model comparison showed that a model with separate learning rates for each recipient best explained behaviour. Young adults were faster to learn when their actions benefitted themselves, compared to helping others. Strikingly, compared to younger adults, older adults showed preserved prosocial learning rates but reduced self-relevant learning rates. Moreover, psychopathic traits were lower in older adults and negatively correlated with prosocial learning. These findings suggest learning how to benefit others is preserved across the lifespan with implications for reinforcement learning and theories of healthy ageing.

Keywords: Prosocial behaviour, ageing, reinforcement learning, computational modelling

Learning associations between actions and their outcomes is fundamental for adaptive behaviour. To date, the majority of studies examining reinforcement learning have tested how we learn associations between actions and outcomes that affect ourselves, and largely focused on these processes in young age, both in humans and other species[1–5]. However, such self-relevant learning may be computationally separable from learning about actions that help other people. Studies suggest slower learning of associations between actions and outcomes when they are about[6] or affect others[7], henceforth referred to as 'prosocial learning'.

Senescence is associated with a multitude of changes including declines in cognitive functioning and perception, but perhaps preservation of affective processing and social cognitive abilities[8–10]. However, less is known about how ageing affects social behaviour, despite the critical importance of this question. Social isolation has been found to be as damaging to physical health as smoking or excessive drinking[11]. Social behaviours that benefit others – prosocial behaviours – are vital for maintaining social bonds and relationships[12] across the lifespan. In addition to the benefits for others, prosociality has been linked to improved life satisfaction[13], mental wellbeing[14], and physical health[15] for the person being prosocial, all of which could contribute to healthy ageing. A key aspect of prosocial behaviour is the ability to learn associations between our own actions and outcomes for other people[7]. Here, we use computational models of reinforcement learning in young and older participants to examine the mechanisms that underpin self-relevant and prosocial learning and associations with healthy individual differences in socio-cognitive ability.

Reinforcement Learning Theory (RLT) provides a powerful framework for understanding and precisely modelling learning[16]. In RLT, prediction errors signal the unexpectedness of outcomes and affect the choices we make in the future. The influence that prediction errors have on choices can be modelled individually through the learning rate, which quantifies the effect of past outcomes on subsequent behaviour. The plausibility of reinforcement learning as a core biological mechanism for action-outcome associations is bolstered by our understanding of neurobiology, with prediction errors encoded by single neurons in the ventral tegmental area[17].

Although essential for successful adaptive behaviour, several studies suggest that our propensity for reinforcement learning declines in later life[10]. Compared to younger adults, older adults show learning impairments particularly when action-outcome associations are probabilistic[18] or reverse[19]. Age-related declines in learning ability have been linked to functional and structural changes in frontostriatal circuits[20,21] and dopamine transmission[18,22], which shows a significant age-related decrease[23–25] and has a key role in coding prediction errors[2,26,27]. Indeed, one study showed that administering L-DOPA, a dopamine precursor, to older adults increased their learning rate[28]. Therefore, if reinforcement learning in general declines in older age, we would hypothesise lower learning rates for both self-relevant and prosocial learning in older, compared to younger, adults.

Alternatively, prosocial learning may depend not only on our learning ability, but also our motivation to help others. Results from experiments using economic games to measure prosociality have found that older adults tend to be more generous[29,30]. There is also evidence of an age-related increase in charitable donations to individuals in need[31]. At work, older adults engage in more prosocial behaviours than younger adults, according to both self-report data and colleagues' ratings[32]. Finally, self-reported altruism and decisions to donate to others have been shown to increase with age[33]. However, one limitation of these studies is that the paradigms often place self and other reward preferences in conflict. Money for the other person depends on less money for oneself. Moreover, older adults generally have higher accumulated wealth, which would be an important confound in studies of monetary exchange[34]. Prosocial learning avoids this confound by separating outcomes for oneself from outcomes for others. If older adults do indeed value outcomes for others more than young adults, we might expect that whilst self-relevant learning declines with ageing, prosocial learning could be preserved. Comparing young and older adults on self-relevant and prosocial learning provides an opportunity to dissociate how possible age-related changes in cognitive ability and social behaviour impact on learning.

While studies point to potential group differences between young and older adults, there is also substantial variability in self and other reward sensitivity. For example,

psychopathy is a key trait associated with decreased prosocial behaviour and altered self and other reward processing[35,36]. Psychopathy has dysfunctional affective-interpersonal features at its core[37,38] but is also characterised by lifestyle and antisocial traits[39]. At the extreme, psychopathy is a severe personality condition linked to poor life outcomes, violence, and criminality[40–42]. However, several studies suggest similar behavioural and neural profiles between community samples with high levels of psychopathic traits[43] and those with clinical diagnoses of psychopathy, consistent with the Research Domains of Criteria (RDoC) conceptualisation of a dimensional approach to psychiatry[44]. This RDoC approach suggests that psychiatric disorders can be thought of as dimensional, rather than categorical, constructs. In a similar vein, psychopathic traits can be captured on a continuum spanning clinical samples and the general population, with a range of scores on that continuum for healthy people. Self-report measures of psychopathic traits that mirror the latent structure of clinical psychopathy measures, comprising antisocial and interpersonal dimensions, are available to use in samples from the general population[39].

Intriguingly, preliminary evidence suggests that ageing may also be associated with changes in psychopathic traits, which could have important implications for our understanding of an ageing population. Epidemiological studies show that criminal activity increases during adolescence then declines in older adulthood[45]. Antisocial and aggressive behaviours also significantly decrease in older age, with young adults (age 16-24 years) having the highest rates of homicide[46]. Even within violent male offenders, psychopathic traits linked to an antisocial lifestyle are negatively correlated with age[47]. In community samples, ageing is associated with a decrease in both the antisocial-lifestyle (antisocial and impulsive behaviours) and affective-interpersonal (lack of empathy and guilt) elements of psychopathic traits[48]. These studies highlight the importance of assessing how differences in psychopathic traits could map on to differences in prosocial behaviours. However, no existing work has examined this question.

Taken together, previous research supports opposing hypotheses for how ageing is associated with self-relevant and prosocial reinforcement learning. On the one hand,

evidence suggests that older adults should be impaired at learning, regardless of the recipient, consistent with ageing-related declines in learning ability and dopamine transmission. On the other hand, potential increases in valuing outcomes for others in older, compared to younger, adults would predict preserved prosocial learning ability but reduced self-relevant learning ability. Finally, we expected variation in psychopathic traits to be associated with learning for others but not self in both age groups.

To distinguish between these competing hypotheses, we tested 75 young (aged 18-36, mean=23.07, 44 females) and 77 older (aged 60-80, mean=69.84, 40 females) adults carefully matched on gender, years of education, and IQ. Participants completed a probabilistic reinforcement-learning task (Figure 1) designed to separate self-relevant (rewards for self) from prosocial learning (rewards for another person), as well as controlling for the general valence of receiving positive outcomes (rewards for neither self or other).

Detailed model comparison showed that a computational model with separate learning rates best explained how people learn associations for different recipients (Figure 2). Young adults were faster to learn when their actions benefitted themselves, compared to when they helped others. Strikingly however, older adults showed a reduced self-bias, with a relative increase in the rate at which they learnt about actions that helped others, compared to themselves (Figure 3a & b). Older adults had significantly reduced levels of psychopathic traits compared to younger adults and in older adults, lower psychopathic trait scores correlated with prosocial learning rates (Figure 4a & b). These effects were not explained by individual differences in IQ, memory or attention abilities. Overall, we show that older adults are less self-biased in reinforcement learning than young adults, and ageing is associated with a decline in psychopathic traits. These findings suggest learning how our actions help others is preserved across the lifespan.

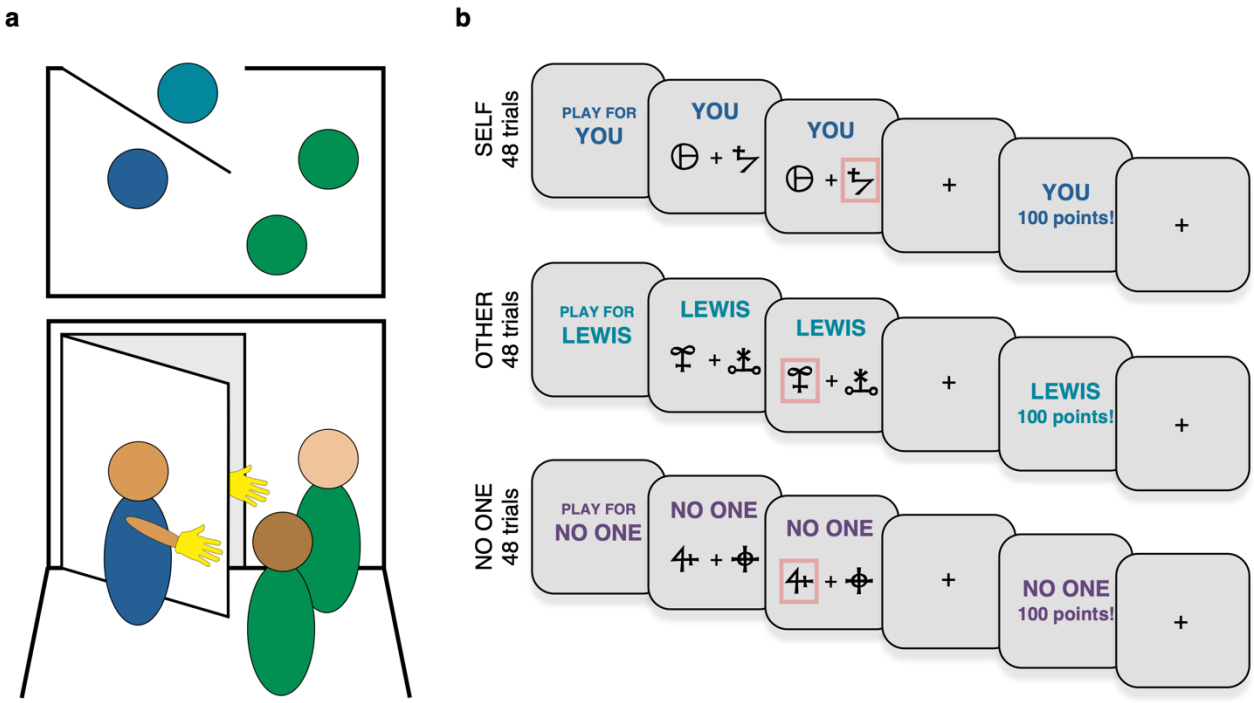

Figure 1. Prosocial learning task and social role assignment. (a) The role assignment procedure involved the participant (dark blue), confederate (light blue) and two experimenters (green). Top: from above showing the positioning of the participant and two experimenters inside the testing room, and the confederate the other side of the door. Bottom: the participant and confederate wore a glove to disguise their identity and waved to each other from either side of the door. Participants were instructed that they would be assigned to roles of 'Player 1' and 'Player 2' but the participant was always assigned to be Player 1. After this procedure participants were informed that they would play a game where they could gain rewards for themselves, the other participant (Player 2) or neither participant. They were told that Player 2 would not play the same game for them and that Player 2 would not know that they may receive an additional bonus based on the choices the participants. This meant that participants' choices were made anonymously and should not be affected by reputational concerns. **(b)** Participants performed a reinforcement learning task ('prosocial learning task') in which they had to learn the probability that abstract symbols were rewarded to gain points. At the beginning of each block, participants were told who they were playing for, either themselves, for the other participant, or in a condition where no one received the outcome. Points from the 'self' condition were converted into additional payment for the participant themselves, points from the 'other' condition were converted into money for Player 2 and points from the 'no one' condition were displayed but not converted into any money for anyone.

Results

We analysed the behaviour of 75 younger adults and 77 older adults who completed the probabilistic reinforcement learning task (Figure 1b), neuropsychological tests, and a measure of psychopathic traits (see Methods). To ensure comparability, older adults with dementia, as diagnosed by the Addenbrooke's Cognitive Examination (ACE)[49], were not included in the study. The two age groups were matched on gender ($\chi^2(1)=0.45$, $p=0.5$) and did not differ in years of education or IQ (Supplementary Table 1). IQ was quantified using age-standardised scores on the Wechsler Test of Adult Reading (WTAR)[50]. We conducted additional analyses controlling for IQ (standardised WTAR score, measured for young and older adults), and memory and attention (memory and attention subscales of the ACE, older adults only). These control analyses showed that our results are not accounted for by general intelligence or executive function (see Methods and Supplementary Information).

Learning occurs for all recipients for both age groups

We first examined whether participants were able to learn for all three recipients to validate their ability to complete the task. We quantified performance as selecting the option associated with a high chance of receiving reward. Participants in both age groups were able to learn to obtain rewards for themselves, another person, and no one. This was demonstrated through average performance above chance level (50%; all ts>15, all ps<0.001) and a significant effect of trial number in predicting trial-by-trial performance (all zs>4.48, ps<0.001) for each separate recipient and age group combination.

Learning rate depends on who receives reward

Next, to quantify learning, we used computational models of reinforcement learning to estimate learning rates (α) and temperature parameters (β), key indices for the speed by which people update their estimates of reward, and the precision with which they make choices, respectively. Models were fitted using a hierarchical approach and

compared using Bayesian model selection as used previously by Huys et al.[51] and Wittmann et al.[52] (see Methods). We tested multiple models that varied with respect to whether learning could be explained by shared or separate free parameters across recipient (self, other, no one). Based on our previous results[7], we examined whether shared or separate learning rates in particular resulted in a better model fit. We used four candidate models:

(i) $1\alpha1\beta$: one α for all three recipients & one β for all three recipients;
(ii) $3\alpha1\beta$: α_{self}, α_{other} & $\alpha_{no\ one}$, one β;
(iii) $2\alpha1\beta$: α_{self} & $\alpha_{not\text{-}self\ [other\ +\ no\ one]}$, one β;
(iv) $3\alpha3\beta$: α_{self}, α_{other}, $\alpha_{no\ one}$, β_{self}, β_{other} & $\beta_{no\ one}$ (see Supplementary Table 2).

Initially, we aimed to establish that both our experimental schedule and our models were constructed in a way that allowed us to disentangle recipient-specific learning rates. To this end, we created synthetic choices using simulations based on each of our four models (see Methods). We fitted the models to the data and assured that the best fitting model was the one that had been used to create the data. In such a way, we established model identifiability, both when considering the exceedance probability (Figure 2a and see Methods) and the number of times a model was identified as the best one (Figure 2b). As a second prerequisite for testing for agent-specific learning rates, we performed parameter recovery using our key model of interest, the $3\alpha1\beta$ model. Over a wide parameter space, we were able to recover the parameters underlying our choice simulation (Figure 2c).

Having established the models were identifiable and parameters recoverable, we performed Bayesian model selection on the data from our participants. Participant's choices were best characterised by the $3\alpha1\beta$ model. This indicated that the learning process underlying the choices is most accurately captured by assuming separate learning rates for each recipient (α_{self}, α_{other} & $\alpha_{no\ one}$). This model fit the data best (exceedance probability=97%; ΔBIC_{int}=122; Figure 2d) and predicted choices well (R^2=51%; see Methods for further details).

229 ***Higher learning rates are associated with better performance***

230

231 As a final check of the robustness of our model and to enable clear interpretation of
232 any differences in learning rate between recipients and age groups (c.f. [53]), we
233 conducted an additional simulation experiment. We simulated data from 10,000
234 participants using the $3\alpha1\beta$ model. This created 30,000 values of α from the three
235 recipient conditions, spanning the full range of possible values from 0 to 1 (see
236 Methods). For each, we quantified the associated performance as the percentage of
237 times the synthetic participant chose the high probability of reward option, averaged
238 across the blocks for the relevant recipient. Plotting the learning rates against
239 performance (Figure 2e) shows that the optimal value of α is approximately 0.55.
240 This is higher than all the values of α found on our task in any recipient condition for
241 either age group. Therefore, higher learning rates were associated with better
242 performance. We further established this link by correlating learning rates and
243 performance in the empirical data from our participants. We found a strong
244 correlation between learning rates and performance overall ($r_{s(150)}$ [95% confidence
245 interval]=0.58 [0.46, 0.68], $p<0.001$; Figure 2f) and in each recipient and age group
246 combination (Supplementary Table 3).

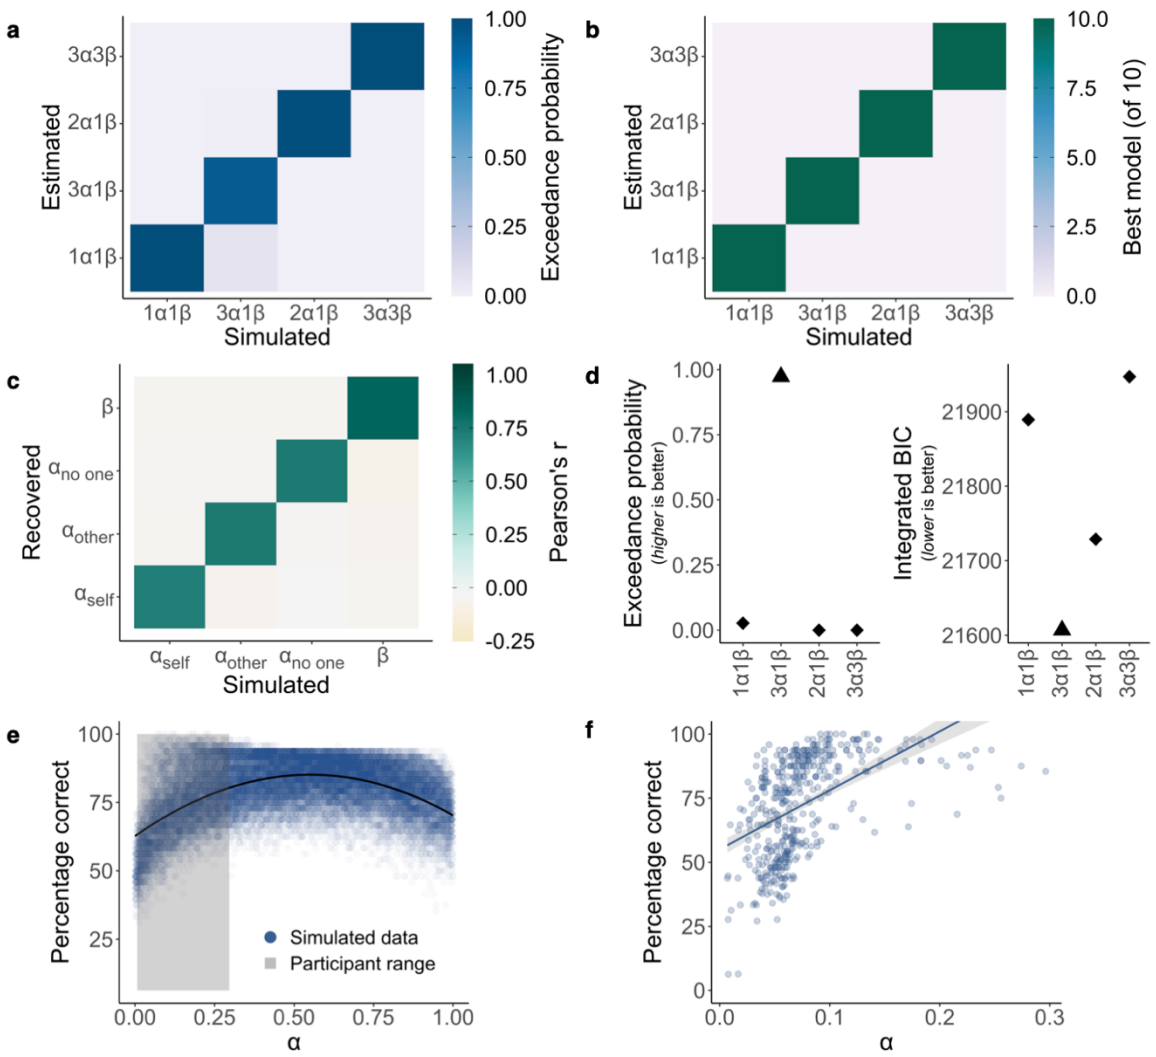

Figure 2. Model identifiability, parameter recovery and simulation. (a) Model identifiability average exceedance probability confusion matrix and **(b)** model identifiability best model selection confusion matrix. Data were simulated from 150 synthetic participants with each of our four models then Bayesian model selection was applied, and this procedure was repeated 10 times. Identifiability is shown by strong diagonals. **(c)** Parameter recovery was performed on data simulated by the winning 3α1β model from 1296 synthetic participants. Confusion matrix represents correlations between simulated and fitted parameters. Strong correlations on the diagonal show parameters can be recovered. **(d)** The 3α1β model (▲) is the best model on both exceedance probability and integrated Bayesian Information Criterion (BIC) fit measures. **(e)** Average percentages of correct choices (high probability of reward option) associated with 30,000 simulated α values (10,000 synthetic participants, 3 recipient conditions) show that an optimal learning rate is approximately 0.55 in this task. The range of α values for our participants was below this peak (grey shading), such that a higher learning rate was associated with better performance. **(f)** Correlation between percentage correct and learning rate across participants. There was a significant correlation between learning rate and accuracy ($r_{s(150)}$=0.58 [0.46, 0.68], p<0.001) (see Supplementary Table 3 for each separate age group and recipient combination; r_s in all cases > 0.46, ps <0.001). Shaded area represents 95% confidence interval.

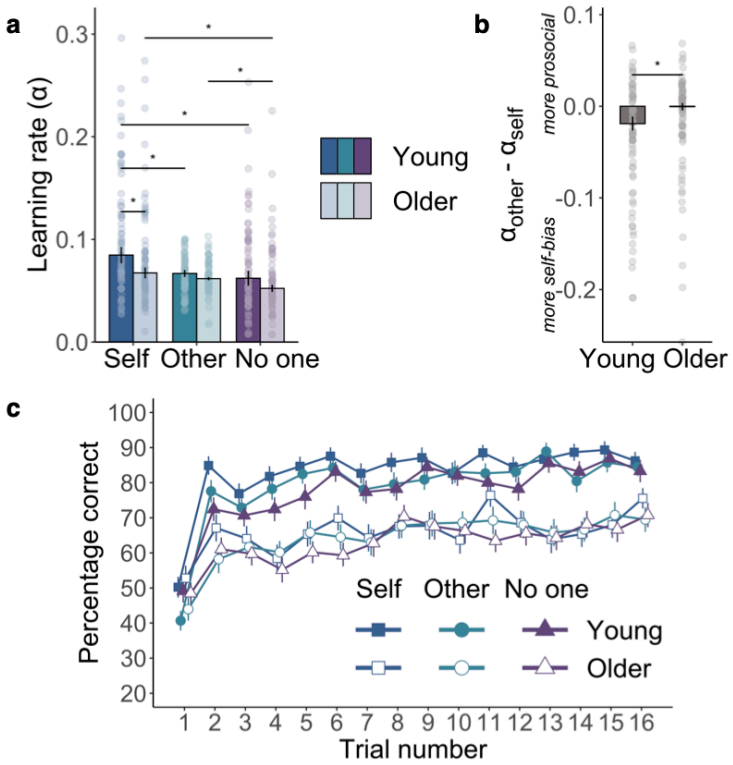

Figure 3. Age-group differences in accuracy and learning rates. (a) Comparison of learning rates from the computational model show older adult's prosocial learning was preserved: learning rates in the other condition did not differ from the self condition or from young adult's prosocial learning rates. Only young participants showed self-bias, n=150 (75 young, 75 older). Bars show group median, error bars are standard error of the median, asterisks represent significant between-group and within-group Wilcoxon t-tests ($p<0.05$). **(b)** Median difference between learning rates in the other and self conditions illustrates the larger self-bias in young, compared to older, adults, n=150 (75 young, 75 older). Error bars are standard error of the median, the asterisk represents the significant age group * recipient [self vs. other] interaction from the robust linear mixed-effects model ($p<0.001$). **(c)** Group-level learning curves showing choice behaviour in the three recipient conditions for each age group. Trials are averaged over the three blocks (48 trials total per recipient presented in three blocks of 16 trials) for the self, other, and no one recipients, n=152 (75 young, 77 older). Points show group mean, error bars are standard error of the mean.

Older adults show a reduced self-bias in learning rates

Next, we used this validated computational model to test our hypotheses as to whether there were group differences in learning rates when learning to reward self, other or no one. Two participants had learning rates for two of the three recipients more than three standard deviations (SDs) above the mean (α_{self} 6.68 & $\alpha_{no\ one}$ 9.64; α_{self} 7.96 & α_{other} 3.78 SDs above the mean) and were excluded from all analysis of learning rates. We analysed the condition-specific learning rates from our best fitting computational model using a robust linear mixed-effects model (RLMM; see Methods). The RLMM fixed effects were age group (young, older), recipient (self, other, no one), as well as the age group * recipient interaction. While this RLMM includes all three recipient conditions, it generates coefficients (main effect and the interaction with age group) contrasting pairs of recipient conditions – [self vs. other] and [self vs. no one]. These are more interpretable than an omnibus test that would not show which recipient conditions were driving an effect or interaction. We followed up these results with planned comparisons, between the older and younger group in each recipient condition, and between pairs of recipient conditions within each age group.

Across age groups, participants showed a higher learning rate when rewards were for themselves, compared to for another person (recipient [self vs. other]: b=-0.024 [-0.034, -0.014], z=-4.79, p<0.001). Importantly however, this pattern differed between age groups. The difference between learning rates for self and other was reduced in older compared to younger adults (recipient [self vs. other] * age group interaction: b=0.016 [0.002, 0.030], z=2.29, p=0.02). Between-group comparisons showed older adults learnt more slowly for themselves compared to younger adults (W=3512, Z=-2.63, $r_{(150)}$=0.22 [0.06, 0.36], p=0.009). However, prosocial learning was preserved, with a Bayes factor suggesting strong evidence of no difference in α_{other} between young and older adults (W=3042, Z=-0.86, $r_{(150)}$=0.07 [0.00, 0.24], p=0.39, BF_{01}=4.26). Within-subject comparisons of α_{self} and α_{other} in each age group showed that young adults had higher learning rates for themselves, relative to another person (V=659, Z=-4.04, $r_{(75)}$=0.47 [0.26, 0.63], p<0.001). In contrast, older adults showed no

312 significant difference between learning rates for self and other (V=1150, Z=-1.45,
313 $r_{(75)}$=0.17 [0.01, 0.38], *p*=0.15, BF_{01}=1.08).
314
315 As expected, across age groups learning was slower for no one than self (recipient
316 [self vs. no one]: *b*=-0.023 [-0.033, -0.013], *z*=-4.57, *p*<0.001). Unlike α_{self} vs. α_{other},
317 learning for self compared to no one did not interact with age group (recipient [self vs.
318 no one] * age group interaction: *b*=0.008 [-0.006, 0.022], *z*=1.15, *p*=0.25). Within-
319 subject comparisons between α_{self} and $\alpha_{no\ one}$ in each age group showed that both
320 groups learnt preferentially for themselves compared to no one (young adults: V=928,
321 Z=-2.62, $r_{(75)}$=0.30 [0.07, 0.51], *p*=0.009; older adults V=901, Z=-2.76, $r_{(75)}$=0.32 [0.09,
322 0.53], *p*=0.006). There was no significant difference between the age groups in $\alpha_{no\ one}$
323 but also no evidence in support of the null (W=3241, Z=-1.61, $r_{(150)}$=0.13 [0.01, 0.29],
324 *p*=0.11, BF_{01}=2.04).
325
326 Considering differences in learning between α_{other} and $\alpha_{no\ one}$, young adults did not
327 differentiate between another person and no one, with strong Bayesian evidence for
328 no difference (V=1533, Z=-0.57, $r_{(75)}$=0.07 [0.00, 0.31], *p*=0.57, BF_{01}=5.08). In
329 contrast, older adults had higher learning rates for another person, compared to no
330 one (V=976, Z=-2.37, $r_{(75)}$=0.27 [0.05, 0.49], *p*=0.02). Crucially, this shows that older
331 adults' lack of differentiation between self and other was not simply because they were
332 insensitive to the recipient condition.
333
334 Finally, we also observed an effect of age on both learning rates overall and
335 temperature parameters. Older adults showed slower learning overall compared to
336 younger adults (*b*=-0.019 [-0.028, -0.009], Z=-3.73, *p*<0.001) and higher levels of
337 exploration of choice options (median β young: 0.05, older: 0.19, W=1511, Z=-4.89,
338 $r_{(150)}$=0.40 [0.26, 0.53], *p*<0.001; Supplementary Figure 1).
339
340 In summary, older adults prosocial learning was preserved at the same rate as young
341 adults, despite age-related declines in self-relevant learning rates. In other words,
342 young adults showed a self-bias in learning, but older adults distinguished between

themselves and others significantly less than the young participants. Only older adults, not young adults, distinguished between rewards for another person and no one.

Participants perform better for themselves, compared to no one

For completeness, we also tested the effects of recipient and age group on trial-by-trial tendency to pick the high reward stimuli (Figure 3c). In addition to the main effect of trial number (b=1.71 [1.29, 2.13], z=7.97, p<0.001), showing learning, these models revealed older adults chose the high reward option less frequently (mean for young: 80%, older: 64%, b=-1.18 [-1.65, -0.70], z=-4.84, p<0.001), and improved less during the task (trial number * age group interaction b=-0.81 [-1.34, -0.27], z=-2.95, p=0.003) across recipients. When averaging across age groups, performance was better for the self (75%), compared to no one (70%; b=-0.36 [-0.05, 0.16], z=-2.28, p=0.02). However, there was not a significant difference between accuracy for other (72%) and self (b=-0.22 [-0.49, 0.05], z=-1.63, p=0.10), or any significant interactions between age group and recipient (bs<0.16, zs<0.96, ps>0.34).

Psychopathic traits are lower in older adults and explain variance in prosocial learning

Finally, we examined individual variability in psychopathic traits, considering age-related differences and influence on prosocial learning. Several studies have suggested that individual differences in psychopathic traits can be meaningfully and accurately captured in community samples and often parallel findings in criminal offenders[43]. Critically, psychopathic traits are closely linked to alterations in social behaviour and willingness to help others. Therefore, we also asked participants to complete the Self-Report Psychopathy Scale (SRP-IV-SF)[39]. The SRP is a measure of psychopathic traits in healthy samples that assesses traits linked to clinical psychopathy, such as antisociality and interpersonal affect. The measure robustly captures the latent structure of clinical psychopathy to enable parallels can be drawn between normal individual differences in the community and clinical samples (see Methods). One participant in each age group had missing questionnaire data and are

not included in these analyses. Psychopathic traits are consistently divided into two components that this scale measures: core affective-interpersonal traits, which capture lack of empathy and guilt; and lifestyle-antisocial traits, which capture impulsivity and antisocial tendencies. Comparing the two age groups on these scales showed that older participants had significantly lower scores than young participants on both the core affective-interpersonal (young mean=24.36, older mean=21.09, W=3558, Z=-3.15, $r_{(148)}$=0.26 [0.11, 0.40], *p*=0.002) and the lifestyle-antisocial subscales (young mean=22.89, older mean=20.27, W=3471, Z=-2.82, $r_{(148)}$=0.23 [0.09, 0.38], *p*=0.005). These findings suggest that both components of psychopathic traits were reduced in older, compared to younger, adults.

Next, we sought to test our hypothesis that individual differences in core psychopathic traits would explain variability in learning rates, specifically for prosocial learning. We observed a significant negative relationship between α_{other} and core psychopathic traits among older participants ($r_{s(74)}$=-0.33 [-0.52, -0.11], *p*=0.005; Figure 4a). Intriguingly, this relationship was significantly more negative (Z=3.28, *p*=0.001) than the equivalent correlation in young adults, which was not significant ($r_{s(74)}$=0.21 [-0.02, 0.42], *p*=0.07; Figure 4b). This pattern of results was the same when correlating 'relative prosocial learning rate' (the difference between α_{other} - α_{self}) with core psychopathic traits. We also conducted control analyses, correlating the same pairs of variables but using partial correlations controlling for β. The negative relationship between prosocial learning (when quantified as α_{other} or α_{other} - α_{self}) and psychopathic traits was still present for older adults, showing that the correlations with learning rates were independent of individual choice exploration (all *p*s<0.05; see Supplementary Table 4). The negative relationship between α_{other} and core psychopathic traits for older adults also remained significant after applying false discovery rate (FDR) correction for multiple comparisons across this correlation and the five other age group & recipient combinations (see Supplementary Table 5). Moreover, the finding that no significant correlations were apparent between psychopathic traits and α_{self} or $\alpha_{no\ one}$ (*p*s>0.15; Supplementary Table 5) suggests a specificity in the relevance of psychopathic traits to prosocial learning in older adults.

Given the age-group differences in levels of psychopathic traits and their correlations with prosocial learning, our final analysis considered whether scores on the core affective-interpersonal psychopathic traits measure mediated the effect of age group on relative prosocial learning rate (α_{other} - α_{self}). A standard mediation model (Figure 5a) did not show evidence for a significant mediation. However, as would be predicted if the link between psychopathic traits and prosocial learning depends on age, we found evidence for a moderated mediation. This revealed core psychopathic traits mediated the effect of age group on relative prosocial learning rate for older adults (unstandardised indirect effect=0.006 [0.001, 0.013], *p*=0.008, proportion mediated=0.32) but not for young adults (unstandardised indirect effect=-0.001 [-0.007, 0.006], *p*=0.65; see Figure 5b for standardised coefficients). In summary, young and older adults differed in levels of psychopathic traits and whether or not psychopathic trait scores explained the extent to which their learning rates were relatively prosocial or self-biased.

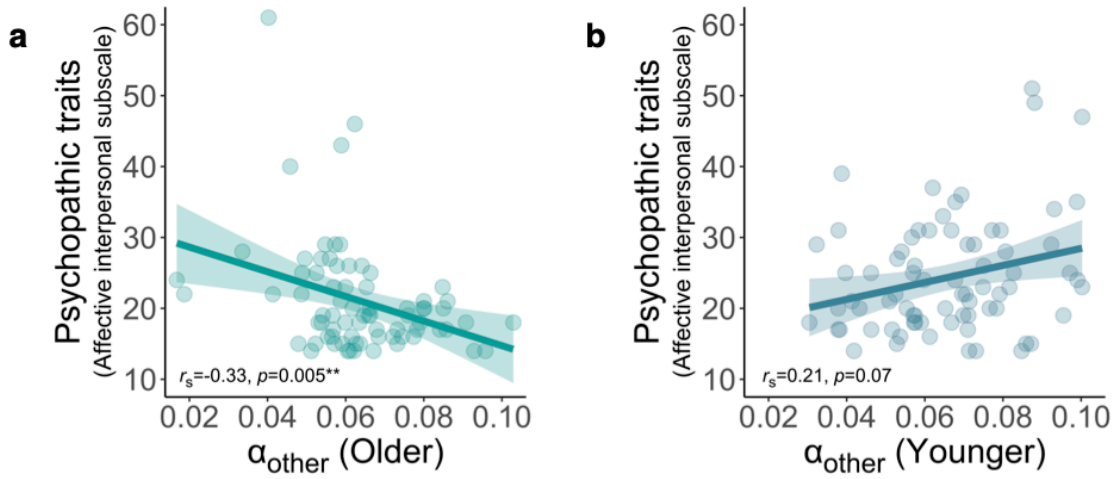

Figure 4. Correlations between prosocial learning rates (α_{other}) and scores on the affective-interpersonal subscale of the Self-Report Psychopathy scale. **(a)** For older adults, levels of psychopathic traits are negatively correlated with prosocial learning rates (r_s=-0.33 [-0.52, -0.11], p=0.005, false discovery rate (FDR) corrected p=0.03). **(b)** There is no significant relationship for young adults (r_s=0.21 [-0.02, 0.42], p=0.07, FDR-corrected p=0.22) and the correlation is significantly more negative (Z=3.28, p=0.001) for older than young adults. This pattern of results is the same when considering correlations between psychopathic traits and the lack of self-bias in learning (α_{other} - α_{self}; not shown). This measure of prosocial learning is also negatively correlated with psychopathic traits in older adults (r_s=-0.25 [-0.45, -0.02], p=0.03) but not younger adults (r_s=0.11 [-0.12, 0.33], p=0.36; difference Z=2.15, p=0.03). Age-group differences in psychopathic traits and the correlation between α_{other} and psychopathic traits, for older adults only, also remained significant when excluding extreme scores (>3 SDs from the mean) on the psychopathic traits measure (Supplementary Tables 6 & 7). Shaded areas represent 95% confidence intervals.

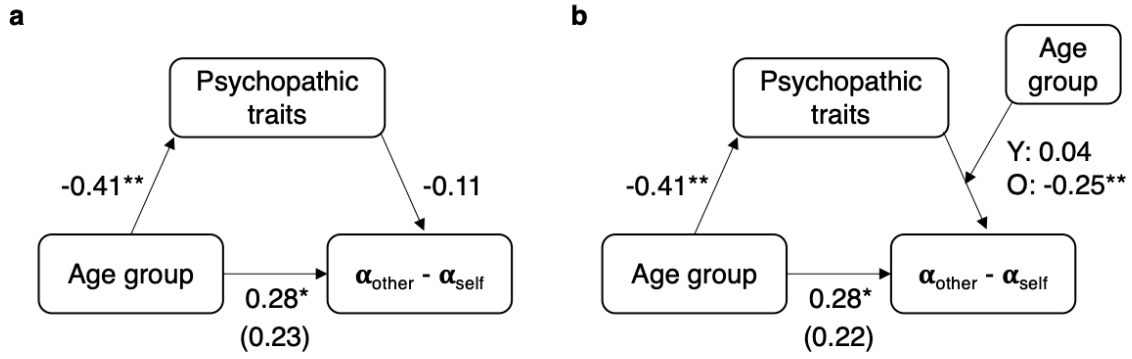

Figure 5. Psychopathic traits mediate the effect of age group on relative prosocial learning (α_{other} - α_{self}) for older adults only. **(a)** A standard mediation model does not show evidence that psychopathic traits mediate the effect of age group on relative prosocial learning. Although accounting for psychopathic traits means the significant direct effect of age group on relative prosocial learning (standardised coefficient=0.28, *p*=0.04) becomes non-significant (standardised coefficient=0.23, *p*=0.11), psychopathic traits do not predict relative prosocial learning overall so there is no mediation. **(b)** Evidence of a moderated mediation is revealed when accounting for differences between young and older adults in how psychopathic traits predict relative prosocial learning. Psychopathic traits are a mediator for older but not younger adults. Psychopathic traits are scores on the affective-interpersonal subscale of the Self-Report Psychopathy scale, Y: young, O: older, asterisks represent significant effects (**p*<0.05, ***p*<0.01).

Discussion

Reinforcement learning is a fundamental process for adaptive behaviour in many species. However, existing studies have largely focused on young people and self-relevant learning, but the decisions we make often occur in a social context[54] and our actions affect outcomes for others. Here, for the first time, we applied computational models of reinforcement learning to the question of ageing-related changes in self-relevant and prosocial learning. We found a clear decrease in learning rates for self-relevant rewards in older, compared to younger, adults. Intriguingly, despite this reduction in self-relevant learning, learning rates for outcomes that affected others did not differ between older and young adults, with Bayesian analyses supporting no difference. Moreover, not only did older adults have preserved prosocial learning rates, age was also associated with lower psychopathic traits, which were specifically linked to prosocial learning ability in older adults.

Models of learning are a powerful tool for understanding prosocial behaviour. By isolating the learning rate, we can precisely examine the influence of reward history on learning. We robustly replicated previous findings that self-relevant learning can be computationally separated from prosocial learning[7], with different learning rates for different recipients providing the best explanation of behaviour. Including separate learning rates improved the model fit and, on average across participants, there was a self-bias, learning rates were higher for self-relevant rewards, compared to when someone else received the reward. However, this self-bias was reduced in older adults who showed higher prosocial learning rates, relative to their own self-relevant learning rates, than young adults. As expected, older adults learnt more slowly for themselves than young adults but in the prosocial condition, the learning rates did not significantly differ between the age groups. Bayesian analyses additionally confirmed that prosocial learning was preserved between young and older adults.

As with much research on age-related changes on cognitive and social tasks, our key finding that self-bias in learning is reduced in older adults could be interpreted as due to changes in ability, motivation, or a combination of these factors. Importantly,

481 learning rates were not associated with executive function or an age-standardised
482 measure of intelligence. We also show that our results remain the same after
483 controlling for these measures. These findings suggest that the observed decline in
484 self-relevant learning rates, but preservation of prosocial learning rates, for older
485 adults are not explained by changes in these broad abilities. Considering learning
486 more specifically, a recent comparison of motivation and ability as explanations for
487 ageing-related reductions in model-based strategies during self-relevant learning
488 supported the limited cognitive abilities account[55]. Our finding that learning rates for
489 self-relevant outcomes were reduced in older adults is in line with a degeneration in
490 the neurocognitive systems required for successful learning. Research combining
491 models of learning with neuroimaging and pharmacological manipulations suggests
492 ageing reduces the ability to generate reward prediction errors[18] (but see[56,57]) due to
493 declines in dopamine functioning[28,58] (also see[59,60]). Differences in motivation could
494 also be applicable for self-relevant learning as the subjective value of financial
495 outcomes is also likely to decrease in older age, due to changes in wealth across the
496 lifespan[34].
497
498 Our findings suggest that despite declines in learning ability associated with ageing,
499 prosocial learning – learning to help others – is preserved. This finding aligns with an
500 emerging literature showing older adults may be more prosocial and less self-biased
501 than younger adults[29,30,61]. The possibility that relatively preserved prosocial learning
502 is related to increased prosocial motivation is in-line with our observed link between
503 learning rates and psychopathic traits. Psychopathic traits were significantly reduced
504 in our older adult sample, dovetailing with similar previous findings on this trait[47,48] and
505 broader trait benevolence[33]. Importantly, we found psychopathic traits in older adults
506 negatively correlated with prosocial learning rates. Self-bias in learning rates (i.e.
507 higher learning rate for self compared to other) was most reduced, and even reversed,
508 in the older people lowest on psychopathic traits. Notably, this negative correlation
509 between psychopathic traits and prosocial learning was only found for older adults.
510 This suggests that age-related differences in prosocial learning could be linked to
511 basic shifts in individual traits and motivations over the lifespan, not just to domain-
512 general reductions in cognitive abilities. The idea that social motivations become more

influential in learning and decision-making with age has also been suggested based on studies of social rewards such as smiling faces or hypothetical time spent with social partners[62–64]. Taken together, this work suggests that strategies to support healthy ageing might benefit from leveraging potentially preserved social motivations.

Importantly, our task included a control condition where no one benefitted. This was essential to establish that the lack of difference between self and other learning rates for older adults was not simply due to an age-related reduction in the absolute dynamic range as maximum learning rates decrease. Older adults had higher learning rates for both others and themselves, compared to this control condition. In contrast, young adults did not differentiate another person from no one. Older adults therefore showed a relative increase in learning rates that was specific to the prosocial condition. This is also evidence against the idea that lower learning rates in older people are reflective of a general reduced sensitivity to who gets the reward. It is interesting to note that the magnitude of the decrease in self-relevant learning rates associated with being older, compared to younger, is similar to the decrease associated with a young person learning for someone else, compared to themselves. The preservation of prosocial learning rates between age groups may seem at odds with the decreased self-relevant learning rates in our sample of older adults and existing evidence of underlying neurobiological deterioration. However, learning from outcomes for self and other have been linked to distinct regions of the brain in humans, shown though neuroimaging[7,65], and causally in monkeys with focal lesions[66].

Taken together, our results add to a growing body of literature suggesting age-related increases in prosocial motivation. If this is the case, the next question is how and why this happens, as there are many possible reasons to be prosocial. For example, prosocial behaviours can be motivated by reputational concerns[67], the ‘warm glow’ of helping[68–70], or vicarious reinforcement from positive outcomes for others[71]. In our procedure, we were very careful to prevent reputational concerns influencing learning to help others. Participants underwent an extensive procedure to introduce them to another participant but to hide information about their age and identity. This meant we could assess tendency for prosocial learning in a situation where reputational

motivations were excluded, and identity-based influences were controlled. Using a reinforcement learning task, in which performance generates positive outcomes for others also focusses on vicarious rewards from the outcome, rather than warm glow associated with the action of helping. Thus, the increase in prosocial learning rates, relative to self-relevant learning rates, suggests older adults are reinforced by outcomes for others and themselves more similarly than younger adults. Many prosocial measures such as the dictator game[72] are costly, requiring direct trade-offs between outcomes for oneself and others. This also makes it hard to detect whether changes are in the value of outcomes for oneself, or outcomes for others, or both. In contrast, separating self-relevant learning, prosocial learning, and the control condition allows us to differentiate increases in the value of prosocial outcomes from decreases in the value of outcomes for the self. Our results are consistent with older adults having both decreased self-relevant learning rates (compared to young adults) and increased prosocial learning rates (compared to their performance in the control condition).

Our results also support the advantages of a model-based approach for understanding both prosocial behaviour and ageing. The model-based parameters were more sensitive to the effects of interest than general measures of performance on the task. The model comparison process is able to provide important information about how the learning process takes place, which cannot be revealed from performance measures alone. We showed that learning was best represented by a separate learning rate for each recipient. Moreover, this approach is able to capture additional latent parameters that drive behaviour, such as the inverse temperature, which indexes how closely participants follow the stimulus value. We demonstrate that a single inverse temperature parameter best explained behaviour during learning across recipient conditions, despite learning rates being distinct.

While our procedure and task have many benefits, it is important to also recognise limitations. To test for age-related differences in prosocial learning, we recruited a group of older adults and a group of younger adults. This increases power to detect differences, but we are unable to assess at what age or how quickly changes take place. Also, while our age groups were matched on years of education and IQ, the

recruitment from university databases or issues around self-selection may mean that the levels of education and IQ in our sample are not completely representative of the general population. Further studies could include samples recruited entirely from the community and participants across the whole adult lifespan. Future studies should also assess the timescale after which older participants reach similar ceiling levels of performance as younger adults, or whether they are never able to reach the same level of performance. Of similar importance is the question of whether older adults may be able to sustain higher levels of motivation over an extended period of time compared to younger adults, possibly compensating for slower learning speeds. In this study, due to time constraints and the presence of several conditions, we are only able to derive conclusions from a limited time window. Moreover, previous research has suggested that individual differences in empathy – the ability to vicariously experience and understand others' affect – might relate to differences in prosocial learning[7]. Empathy is positively associated with affective-interpersonal psychopathic traits and might also relate to motivation to help others[73,74]. Further studies could also assess how empathy predicts changes in prosocial learning across the lifespan. Finally, future studies could manipulate the identity of the recipient as we show a preservation in tendency to help others when there are no particular characteristics known about the other person, but these effects might additionally be modulated by factors such as perceived social distance.

To conclude, we find new evidence that despite declines in self-relevant learning in older adults, the ability to learn which actions benefit others is preserved. Moreover, the bias with which people favour self-relevant outcomes is reduced. Not only do older adults have relatively preserved prosocial learning, they also report lower levels of core psychopathic traits that index lack of empathy and guilt, and this trait difference is linked to the changes in prosocial learning. These findings could have important implications for our understanding of reinforcement learning and theoretical accounts of healthy ageing.

Materials and Methods

Participants

We recruited 80 young participants and 80 older participants using the same recruitment methods in order to match the samples as closely as possible. Participants were recruited from university databases, which included students and members of the community, social media, and adverts in local newspapers. We excluded anyone who was currently studying or had previously studied psychology and no one took part for course credit. Additional exclusion criteria were previous or current neurological or psychiatric disorder, non-normal or non-corrected to normal vision and, for the older sample, scores on the Addenbrooke's Cognitive Examination that indicate potential dementia (cut-off score 82)[49]. This sample size gave us 88% power to detect a medium size effect (d=0.5).

Five young and three older participants were excluded due to: diagnosis of a psychiatric disorder at the time of testing (1 young participant); previous study of psychology (2 young participants); and incomplete or low-quality data (2 young and 3 older participants). This left a final sample of 152 participants, 75 young adults (44 females aged 18-36, mean=23.07) and 77 older adults (40 females aged 60-80, mean=69.84). Two older participants were excluded from all analyses involving learning rates due to each having two learning rate estimates more than three standard deviations (SDs) above the mean (for one α_{self} was 6.68 SDs above the mean & α_{no} one 9.64 SDs above the mean; for the second α_{self} was 7.96 & α_{other} 3.78 SDs above the mean). One further participant from each age group was missing data on the SRP measure so are not included in the relevant analyses.

Participants were paid at a rate of £10 per hour plus an additional payment of up to £5 depending on the number of points they earned for themselves during the task. They were also told the number of points that they earned in the prosocial condition would translate into an additional payment of up to £5 for the other participant (see details of the task below). All participants provided written informed consent and the study was

640 approved by the Oxford University Medical Sciences Inter Divisional Research Ethics
641 Committee and National Health Service Ethics.
642
643 ***Prosocial learning task***
644
645 The prosocial learning task is a probabilistic reinforcement learning task, with rewards
646 in one of three recipient conditions: for the participant themselves (*self*), for another
647 participant (*other*; prosocial condition), and for *no one* (control condition). Each trial
648 presents two symbols, one associated with a high (75%) probability of gaining points
649 and the other with a low (25%) probability of gaining points. The two symbols were
650 randomly assigned to the left or right side of the screen and selected via a
651 corresponding button press. Participants select a symbol then receive feedback on
652 whether they obtained points or not (see Figure 1b) so learn over time which symbol
653 maximises rewards. The experiment was subdivided into blocks, i.e. 16 trials pairing
654 the same two symbols for the same recipient. Participants completed three blocks, a
655 total of 48 trials, in each recipient condition, resulting in 144 trials overall (see
656 Supplementary Information for trial structure). Blocks for different recipients were
657 pseudo-randomly ordered such that the same recipient block did not occur twice in a
658 row.
659
660 On trials in the *self* condition, points translated into increased payment for the
661 participant themselves. These blocks started with “play for you” displayed and had the
662 word “you” at the top of each screen. Blocks in the *no one* condition had “no one” in
663 place of “you” and points were not converted into any extra payment for anyone. In
664 the prosocial *‘other’* condition, participants earned points that translated into additional
665 payment for a second participant, actually a confederate. Participants were told that
666 this payment would be given anonymously, they would never meet the other person,
667 and that the person was not even aware of them completing this task (see
668 Supplementary Information). The name of the confederate, gender-matched to the
669 participant, was displayed on these blocks at the start and on each screen (Figure 1b).
670 Thus, participants were explicitly aware who their decisions affected on each trial.

671 Stimuli were presented using Presentation (Neurobehavioral Systems –
672 https://www.neurobs.com/).
673
674 ***Questionnaire measures***
675
676 **Dementia screening and executive function.** Older adults were screened for
677 dementia using the Addenbrooke's Cognitive Examination (ACE-III)[49]. The ACE
678 examines five cognitive domains; attention, memory, language, fluency and
679 visuospatial abilities. The ACE-III is scored out of 100 and as a screening tool, a cut-
680 off score of 82/100 denotes significant cognitive impairment. We also used scores on
681 the attention and memory domains in control analyses as proxies for executive
682 function in older adults.
683
684 **General intelligence.** All participants completed the Wechsler Test of Adult Reading
685 (WTAR)[50] as a measure of IQ. The WTAR requires participants to pronounce 50 words
686 that have unusual grapheme-to-phoneme translation. This means the test measures
687 reading recognition or existing knowledge of the words, rather than ability to apply
688 rules for pronunciation. The WTAR was developed and standardised with the
689 Wechsler Memory and Adult Intelligence Scales and correlates highly with these
690 measures[75]. Standardisation involves adjusting for healthy age-related differences.
691 The test is suitable for participants aged 16-89, covering our full sample, and scores
692 in older age have been shown to correlate with cognitive ability earlier in life[76].
693
694 **Psychopathic traits.** Participants completed the short form of the Self-Report
695 Psychopathy Scale 4th Edition (SRP-IV-SF)[39]. This scale consists of 29 items, 7 each
696 measuring: interpersonal, affective, lifestyle and antisocial tendencies (plus 'I have
697 been convicted of a serious crime'). We used the two-factor grouping, summing the
698 core, affective-interpersonal items and separately, the lifestyle-antisocial items for use
699 in analysis. The robust psychometric properties of this measure have been established
700 in community[77] and offender populations through construct and convergent validity[78],
701 internal consistency, and reliability[39].
702

703 ***Procedure***

704

705 **Role assignment.** To enhance belief that points earned in the prosocial condition
706 benefitted another person, we conducted a role assignment procedure based on a set-
707 up used in several studies of social decision-making[79,80]. Participants were instructed
708 not to speak and wore a glove to hide their identity. A second experimenter brought
709 the confederate, also wearing a glove, to the other side of the door. Participants only
710 ever saw the gloved hand of the confederate, but they waved to each other to make it
711 clear there was another person there (Figure 1a). The experimenter tossed a coin to
712 determine who picked a ball from the box first and then told the participants which
713 roles they had been assigned to, based on the ball they picked. Our procedure
714 ensured that participants always ended up in the role of the person performing the
715 prosocial learning experiment. Participants were unaware of the age of the other
716 person, but the experimenter used a name for them suggesting their gender was the
717 same as the participant.

718

719 **Task procedure.** Participants received instructions for the learning task and how the
720 points they earned would be converted into money for themselves and for the other
721 participant. Instructions included that the two symbols were different in how likely it
722 was that choosing them lead to points but that which side they appeared on the screen
723 was irrelevant. Participants then completed one block of practice trials before the main
724 task and were aware outcomes during the practice did not affect payment for anyone.
725 After the task, participants completed the measure of psychopathic traits and the
726 dementia screening

727 ***Computational modelling***

728

729 We modelled learning during the task with a reinforcement learning algorithm[16],
730 creating variations of the models through the number of parameters used to explain
731 the learning rate and temperature parameters in the task[81]. The basis of the
732 reinforcement learning algorithm is the expectation that an action (or stimulus) a will
733 provide reward on the following trial. This expected value, $Q_{t+1}(a)$ is quantified as a
734 function of current expectations $Q_t(a)$ and the prediction error δ_t, which is scaled by
735 the learning rate α:

736

737 $$Q_{t+1}(a) = Q_t(a) + \alpha \times \underbrace{[r_t - Q_t(a)]}_{\text{Prediction error } \delta_t}$$

738

739 Where δ_t, the prediction error, is the difference between the actual reward experienced
740 on the current trial r_t (1 for reward and 0 for no reward) minus the expected reward on
741 the current trial $Q_t(a)$.

742

743 The learning rate α therefore determines the influence of the prediction error. A low
744 learning rate means new information affects expected value to a lesser extent. The
745 softmax link function quantifies the relationship between the expected value of the
746 action $Q_t(a)$ and the probability of choosing that action on trial t:

747

748 $$p_t[(a|Q_t(a))] = \frac{e^{(Q_t(a)/\beta)}}{\sum_{a'} e^{(Q_t(a')/\beta)}}$$

749

750 The temperature parameter β represents the noisiness of decisions – whether the
751 participant explores or always chooses the option with the highest expected value. A
752 high value for β means choices seem random as they are equally likely irrespective of
753 the expected value of each option. A low β leads to choosing the option with the
754 greatest expected value on all trials.

755

756

757 ***Model fitting***

758

759 We used MATLAB 2019b (The MathWorks Inc) for all model fitting and comparison.
760 To fit the variations of the learning model (see below) to (real and simulated)
761 participant data we used an iterative maximum a posteriori (MAP) approach previously
762 described[51,52]. This method provides a better estimation than a single-step maximum
763 likelihood estimation (MLE) alone by being less susceptible to the influence of outliers.
764 It does this via implementing two levels: the lower level of the individual subjects and
765 the higher-level reflecting our full sample. For the real participant data, we fit the model
766 across groups to provide the most conservative comparison, so this full sample
767 combined young and older participants.

768

769 For the MAP procedure, we initialized group-level Gaussians as uninformative priors
770 with means of 0.1 (plus some added noise) and variance of 100. During the
771 expectation, we estimated the model parameters (α and β) for each participant using
772 an MLE approach calculating the log-likelihood of the subject's series of choices given
773 the model. We then computed the maximum posterior probability estimate, given the
774 observed choices and given the prior computed from the group-level Gaussian, and
775 recomputed the Gaussian distribution over parameters during the maximisation step.
776 We repeated expectation and maximization steps iteratively until convergence of the
777 posterior likelihood summed over the group, or a maximum of 800 steps. Convergence
778 was defined as a change in posterior likelihood <0.001 from one iteration to the next.
779 Note that bounded free parameters were transformed from the Gaussian space into
780 the native model space via appropriate link functions (e.g. a sigmoid function in the
781 case of the learning rates) to ensure accurate parameter estimation near the bounds.
782 The detailed code for the models and implementation of the fitting algorithm can be
783 found here: https://osf.io/xgw7h/?view_only=bea7b82435b344c5a26f80fd21d8ce19.

784

785 ***Model comparison***

786

787 Our hypotheses generated four models to compare that differed in whether the model
788 parameters (α and β) for each participant had one value across recipient conditions or

789 depended on the recipient (self, other and no one). For model comparison, we
790 calculated the Laplace approximation of the log model evidence (more positive values
791 indicating better model fit[82]) and submitted these to a random-effects analysis using
792 the spm_BMS routine[83] from SPM 8 (http://www.fil.ion.ucl.ac.uk/spm/software/spm8/).
793 This generates the exceedance probability: the posterior probability that each model
794 is the most likely of the model set in the population (higher is better, over .95 indicates
795 strong evidence in favour of a model). For the models of real participant data, we also
796 calculated the integrated BIC (lower is better[51,52]) and R^2 as additional measures of
797 model fit. To calculate the model R^2, we extracted the choice probabilities generated
798 for each participant on each trial from the winning model. We then took the squared
799 median choice probability across participants. The $3\alpha1\beta$ model had the best evidence
800 on all measures (see Supplementary Table 2).
801
802 ***Simulation experiments***
803
804 We used simulation experiments to assess that our experiment allowed us to
805 dissociate models of interest, as well as parameters of interest within the winning
806 model. We simulated data from all four models to establish that our model comparison
807 procedure (see above) could accurately identify the best model across a wide range
808 of parameter values. For this model identifiability analysis, we simulated data from 150
809 participants, drawing parameters from distributions commonly used in the
810 reinforcement learning literature[84,85]. Learning rates (α) were drawn from a beta
811 distribution (betapdf(parameter,1.1,1.1)) and softmax temperature parameters (β)
812 from a gamma distribution (gampdf(parameter,1.2,5)). We fitted the models to this
813 simulated data set using the same MAP process as applied to the experimental
814 participants' data and repeated this whole procedure 10 times. By plotting the
815 confusion matrices of average exceedance probability (across the 10 runs; Figure 2a)
816 and how many times each model won (Figure 2b), we show the models are identifiable
817 using our model comparison process.
818
819 Our winning model contained 4 free parameters (α_{self}, α_{other}, $\alpha_{no\ one}$, β). To assess the
820 reliability of this model and the interpretability of the free parameters, we also

performed parameter recovery on simulated data (see Supplementary Information for procedure) as recommended for modelling analyses that use a 'data first' approach[81,86]. We simulated choices 1296 times using our experimental schedule and fitted them using MAP. We found strong Pearson's correlations between the true simulated and fitted parameter values (all *r*s>0.7, see Figure 2c), suggesting our experiment was well suited to estimate the model's parameters.

Finally, we conducted a principled simulation experiment to identify the optimal learning rate in our task and examine the link between learning rates and performance. We simulated data from 10,000 participants with learning rates (α) and softmax temperature parameters (β) drawn from beta and gamma distributions respectively, as described above. We bounded the β parameters at 0 and 0.3 to reflect the range shown by our participants. In line with our winning model, simulated participants had a separate learning rate for each recipient condition, and these spanned the full range of possible α values from 0 to 1. This generated 30,000 learning rate values (10,000 participants and three conditions). For each we quantified performance as the proportion of times the simulated participant chose the 'correct' high reward option in the relevant condition. Results showed that average performance improved as learning rates increased, up to learning rates of approximately 0.55 (Figure 2e). Crucially, this optimal alpha was above the highest learning rate shown by any of our participants in any recipient condition, meaning a higher learning rate on our task was associated with better performance.

Statistical analysis

Analysis of group and recipient differences in the fitted model parameters and behavioural data was run in R[87] with R Studio[88]. We used a robust linear mixed-effect model (RLMM; rlmer function; robustlmm package[89] to predict learning rates and generalised linear mixed-effects models (glmer function; lme4 package[90]) for the trial-by-trial data (binary outcome of choosing the high vs. low reward option). We used (robust / generalised) linear mixed-effects models as these account for the within-subject nature of the recipient manipulation and do not rely on parametric

assumptions. Additionally, unlike an anova model or omnibus test, an RLMM generates coefficients (with confidence intervals and significance values) for terms that compare pairs of factor levels (for example self vs. other). This approach is more informative than an anova when the factor has more than two levels, which is our case as the recipient factor has three levels (self, other, no one). Each linear mixed-effects model had fixed effects of age group, recipient (self, other, no one), and their interaction, plus a random subject-level intercept. Analysis of trial-by-trial choices also included trial number in the fixed terms, interacting with recipient and group (including the three-way interaction), and in the random terms, interacting with recipient. In the analysis of learning rates controlling for IQ, standardised scores on the WTAR were also included as a fixed term (Supplementary Table 8). Correlations of learning rates with psychopathic traits (Supplementary Table 4 & 5) and neuropsychological measures (Supplementary Table 9) were calculated with Spearman's Rho nonparametric tests. To control for IQ and executive function in the associations between older adults' prosocial learning rates and psychopathic traits, we ran partial correlations each controlling for one of WTAR, ACE memory and ACE attention scores (Supplementary Table 10).

For simple and post hoc comparisons, we used two-sided nonparametric tests as outcome variables violated normality assumptions. Effect sizes and confidence intervals for paired and independent nonparametric comparisons were calculated using the *cohens_d* and *wilcox_effsize* functions respectively from the rstatix package[91]. Bayes factors (BF_{01}) for non-significant results were calculated using nonparametric paired and independent t-tests in JASP[92] with the default prior. BF_{01} corresponds to how many times more likely the data are under the null hypothesis of no difference than under the alternative hypothesis that there is a difference. A BF_{01} larger than 3 (equal to BF_{10} less than 1/3) is considered substantial evidence in favour of the null hypothesis whereas a BF_{01} between 1/3 and 3 indicates the data cannot clearly differentiate between hypotheses[93]. Median learning rates and their standard errors for plotting were calculated using bootstrapping with 1,000 samples.

884 For the mediation analysis, we used the *mediate* function (mediation package[94])
885 combined with robust linear models (*rlm* function, MASS package[95]). This method
886 estimates the unstandardised indirect effects, with 95% confidence intervals and
887 significance, through a bootstrapping procedure with 10,000 bootstrapped samples.
888 The outcome in the mediation model was relative prosocial learning rate (α_{other} -
889 α_{self}), the predictor was age group, and the mediator was psychopathic traits on the
890 core affective-interpersonal subscale of the SRP-IV-SF. We calculated two mediation
891 models. The first, a standard mediation model, included an indirect path from age
892 group to relative prosocial learning rate via psychopathic traits as a mediator (Figure
893 5a). The second included an interaction between age group and psychopathic traits
894 in predicting relative prosocial learning rate to allow the possibility of a moderated
895 mediation[96]. Specifically, this type of moderated mediation examines whether the
896 effect of the mediator (in this case psychopathic traits) on the outcome (α_{other} - α_{self})
897 is moderated by the predictor (age group)[97,98]. In other words, psychopathic traits
898 could be a mediator for one age group but not the other (Figure 5b).

Data availability

Data are available at:
https://osf.io/xgw7h/?view_only=bea7b82435b344c5a26f80fd21d8ce19.

Code availability

Code for modelling and analysis is available at:
https://osf.io/xgw7h/?view_only=bea7b82435b344c5a26f80fd21d8ce19.

Acknowledgements

This work was supported by a Medical Research Council Fellowship (MR/P014097/1), a Christ Church Junior Research Fellowship, and a Christ Church Research Centre Grant to PL; a Wellcome Trust Principal Fellowship to MH; NIHR Biomedical Research Centre, Oxford. The Wellcome Centre for Integrative Neuroimaging is supported by core funding from the Wellcome Trust (203139/Z/16/Z).

We are grateful to Craig Neumann for assistance with the Self-Report Psychopathy Scale. We are also grateful to the many people who acted as confederates for us during the study.

Author contributions

PL designed the study. AA, LH, DD & PL collected the data. JC, MW & PL analysed the data. JC, MW, MH & PL wrote the paper.

Competing interests

The authors declare no competing interests

References

1. Cohen, M. X. Neurocomputational mechanisms of reinforcement-guided learning in humans: A review. *Cogn. Affect. Behav. Neurosci.* **8**, 113–125 (2008).
2. Glimcher, P. W. Understanding dopamine and reinforcement learning: The dopamine reward prediction error hypothesis. *Proc. Natl. Acad. Sci. U. S. A.* **108**, 15647–15654 (2011).
3. Niv, Y. Reinforcement learning in the brain. *J. Math. Psychol.* **53**, 139–154 (2009).
4. O'Doherty, J. P., Lee, S. W. & McNamee, D. The structure of reinforcement-learning mechanisms in the human brain. *Curr. Opin. Behav. Sci.* **1**, 94–100 (2015).
5. Schultz, W. Neuronal Reward and Decision Signals: From Theories to Data. *Physiol. Rev.* **95**, 853–951 (2015).
6. Lockwood, P. L. *et al.* Neural mechanisms for learning self and other ownership. *Nat. Commun.* **9**, 4747 (2018).
7. Lockwood, P. L., Apps, M. A. J., Valton, V., Viding, E. & Roiser, J. P. Neurocomputational mechanisms of prosocial learning and links to empathy. *Proc. Natl. Acad. Sci.* **113**, 201603198–201603198 (2016).
8. Blanchard-Fields, F. Everyday Problem-solving and Emotion: An Adult Developmental Perspective. *Psychol. Sci.* 26–31 (2007).
9. Blanchard-Fields, F. & Hess, T. M. The Social Cognitive Perspective and the Study of Aging. in *Social Cognition and Aging* (eds. Hess, T. M. & Blanchard-Fields, F.) 1–14 (Academic Press, 1999). doi:10.1016/B978-012345260-3/50002-1.

954 10. Samanez-Larkin, G. R. & Knutson, B. Decision making in the ageing brain:
955 changes in affective and motivational circuits. *Nat. Rev. Neurosci.* **16**, 278–289
956 (2015).
957 11. Holt-Lunstad, J., Smith, T. B., Baker, M., Harris, T. & Stephenson, D. Loneliness
958 and social isolation as risk factors for mortality: a meta-analytic review. *Perspect.*
959 *Psychol. Sci. J. Assoc. Psychol. Sci.* **10**, 227–237 (2015).
960 12. Fehr, E. & Fischbacher, U. The nature of human altruism. *Nature* **425**, 785–791
961 (2003).
962 13. Buchanan, K. E. & Bardi, A. Acts of Kindness and Acts of Novelty Affect Life
963 Satisfaction. *J. Soc. Psychol.* **150**, 235–237 (2010).
964 14. Raposa, E. B., Laws, H. B. & Ansell, E. B. Prosocial Behavior Mitigates the
965 Negative Effects of Stress in Everyday Life. *Clin. Psychol. Sci.* **4**, 691–698
966 (2016).
967 15. Post, S. G. Altruism, Happiness, and Health: It's Good to Be Good. *Int. J. Behav.*
968 *Med.* **12**, 66–77 (2005).
969 16. Sutton, R. S. & Barto, A. G. *Reinforcement learning: An introduction*. (MIT Press,
970 2011).
971 17. Schultz, W. Multiple dopamine functions at different time courses. *Annu. Rev.*
972 *Neurosci.* **30**, 259–288 (2007).
973 18. Eppinger, B., Hämmerer, D. & Li, S.-C. Neuromodulation of reward-based
974 learning and decision making in human aging. *Ann. N. Y. Acad. Sci.* **1235**, 1–17
975 (2011).
976 19. Mell, T. *et al.* Effect of aging on stimulus-reward association learning.
977 *Neuropsychologia* **43**, 554–563 (2005).

978 20. Samanez-Larkin, G. R., Levens, S. M., Perry, L. M., Dougherty, R. F. & Knutson,
979 B. Frontostriatal white matter integrity mediates adult age differences in
980 probabilistic reward learning. *J. Neurosci. Off. J. Soc. Neurosci.* **32**, 5333–5337
981 (2012).
982 21. Samanez-Larkin, G. R., Worthy, D. A., Mata, R., McClure, S. M. & Knutson, B.
983 Adult age differences in frontostriatal representation of prediction error but not
984 reward outcome. *Cogn. Affect. Behav. Neurosci.* **14**, 672–682 (2014).
985 22. Hammerer, D. & Eppinger, B. Dopaminergic and Prefrontal Contributions to
986 Reward-Based Learning and Outcome Monitoring During Child Development and
987 Aging. *Dev. Psychol.* **48**, 826–874 (2012).
988 23. Bäckman, L., Nyberg, L., Lindenberger, U., Li, S.-C. & Farde, L. The correlative
989 triad among aging, dopamine, and cognition: Current status and future prospects.
990 *Neurosci. Biobehav. Rev.* **30**, 791–807 (2006).
991 24. Dreher, J.-C., Meyer-Lindenberg, A., Kohn, P. & Berman, K. F. Age-related
992 changes in midbrain dopaminergic regulation of the human reward system. *Proc.*
993 *Natl. Acad. Sci.* **105**, 15106–15111 (2008).
994 25. Li, S.-C., Lindenberger, U. & Bäckman, L. Dopaminergic modulation of cognition
995 across the life span. *Neurosci. Biobehav. Rev.* **34**, 625–630 (2010).
996 26. Schultz, W. Updating dopamine reward signals. *Curr. Opin. Neurobiol.* **23**, 229–
997 238 (2013).
998 27. Schultz, W. Dopamine reward prediction-error signalling: a two-component
999 response. *Nat. Rev. Neurosci.* **17**, 183–195 (2016).
1000 28. Chowdhury, R. *et al.* Dopamine restores reward prediction errors in old age. *Nat.*
1001 *Neurosci.* **16**, 648–653 (2013).

1002 29. Engel, C. Dictator games: a meta study. *Exp. Econ.* **14**, 583–610 (2011).
1003 30. Matsumoto, Y., Yamagishi, T., Li, Y. & Kiyonari, T. Prosocial Behavior Increases
1004 with Age across Five Economic Games. *PLOS ONE* **11**, e0158671 (2016).
1005 31. Sze, J. A., Gyurak, A., Goodkind, M. S. & Levenson, R. W. Greater Emotional
1006 Empathy and Prosocial Behavior in Late Life. *Emotion* **12**, 1129–1140 (2012).
1007 32. Ng, T. W. H. & Feldman, D. C. The relationship of age to ten dimensions of job
1008 performance. *J. Appl. Psychol.* **93**, 392–423 (2008).
1009 33. Hubbard, J., Harbaugh, W. T., Srivastava, S., Degras, D. & Mayr, U. A general
1010 benevolence dimension that links neural, psychological, economic, and life-span
1011 data on altruistic tendencies. *J. Exp. Psychol. Gen.* **145**, 1351–1358 (2016).
1012 34. Cheung, F. & Lucas, R. E. When Does Money Matter Most? Examining the
1013 Association between Income and Life Satisfaction over the Life Course. *Psychol.*
1014 *Aging* **30**, 120–135 (2015).
1015 35. Buckholtz, J. W. *et al.* Mesolimbic Dopamine Reward System Hypersensitivity in
1016 Individuals with Psychopathic Traits. *Nat. Neurosci.* **13**, 419–421 (2010).
1017 36. White, B. A. Who cares when nobody is watching? Psychopathic traits and
1018 empathy in prosocial behaviors. *Personal. Individ. Differ.* **56**, 116–121 (2014).
1019 37. White, S. F. *et al.* Reduced Amygdala Response in Youths With Disruptive
1020 Behavior Disorders and Psychopathic Traits: Decreased Emotional Response
1021 Versus Increased Top-Down Attention to Nonemotional Features. *Am. J.*
1022 *Psychiatry* **169**, 750–758 (2012).
1023 38. Marsh, A. A. *et al.* Empathic responsiveness in amygdala and anterior cingulate
1024 cortex in youths with psychopathic traits. *J. Child Psychol. Psychiatry* **54**, 900–
1025 910 (2013).

1026 39. Paulhus, D. L., Neumann, C. S., Hare, R. D., Williams, K. M. & Hemphill, J. F.
1027 *Self-report Psychopathy Scale 4th Edition (SRP 4) Manual*. (MHS, Multi-Health
1028 Systems Incorporated, 2017).
1029 40. Leistico, A.-M. R., Salekin, R. T., DeCoster, J. & Rogers, R. A large-scale meta-
1030 analysis relating the Hare measures of psychopathy to antisocial conduct. *Law*
1031 *Hum. Behav.* **32**, 28–45 (2008).
1032 41. Blais, J., Solodukhin, E. & Forth, A. E. A Meta-Analysis Exploring the
1033 Relationship Between Psychopathy and Instrumental Versus Reactive Violence.
1034 *Crim. Justice Behav.* **41**, 797–821 (2014).
1035 42. Asscher, J. J. *et al.* The relationship between juvenile psychopathic traits,
1036 delinquency and (violent) recidivism: A meta-analysis. *J. Child Psychol.*
1037 *Psychiatry* **52**, 1134–1143 (2011).
1038 43. Seara-Cardoso, A., Neumann, C., Roiser, J., McCrory, E. & Viding, E.
1039 Investigating associations between empathy, morality and psychopathic
1040 personality traits in the general population. *Personal. Individ. Differ.* **52**, 67–71
1041 (2012).
1042 44. Insel, T. R. The NIMH Research Domain Criteria (RDoC) Project: Precision
1043 Medicine for Psychiatry. *Am. J. Psychiatry* **171**, 395–397 (2014).
1044 45. Liberman, A. *The long view of crime: A synthesis of longitudinal research*.
1045 (Springer, 2008).
1046 46. Homicide in England and Wales - Office for National Statistics.
1047 https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/articles/
1048 homicideinenglandandwales/yearendingmarch2018#what-do-we-know-about-
1049 suspects.

1050 47. Huchzermeier, C. *et al.* Are there age-related effects in antisocial personality
1051 disorders and psychopathy? *J. Forensic Leg. Med.* **15**, 213–218 (2008).
1052 48. Gill, D. J. & Crino, R. D. The Relationship between Psychopathy and Age in a
1053 Non-Clinical Community Convenience Sample. *Psychiatry Psychol. Law* **19**,
1054 547–557 (2012).
1055 49. Hsieh, S., Schubert, S., Hoon, C., Mioshi, E. & Hodges, J. R. Validation of the
1056 Addenbrooke's Cognitive Examination III in frontotemporal dementia and
1057 Alzheimer's disease. *Dement. Geriatr. Cogn. Disord.* **36**, 242–250 (2013).
1058 50. The Psychological Corporation. *Wechsler test of adult reading*. (Harcourt
1059 Assessment, 2001).
1060 51. Huys, Q. J. *et al.* Disentangling the roles of approach, activation and valence in
1061 instrumental and pavlovian responding. *PLoS Comput Biol* **7**, e1002028 (2011).
1062 52. Wittmann, M. K. *et al.* Global reward state affects learning and activity in raphe
1063 nucleus and anterior insula in monkeys. *Nat. Commun.* **11**, 3771 (2020).
1064 53. Zhang, L., Lengersdorff, L., Mikus, N., Gläscher, J. & Lamm, C. Using
1065 reinforcement learning models in social neuroscience: frameworks, pitfalls and
1066 suggestions of best practices. *Soc. Cogn. Affect. Neurosci.* **15**, 695–707 (2020).
1067 54. Olsson, A., Knapska, E. & Lindström, B. The neural and computational systems
1068 of social learning. *Nat. Rev. Neurosci.* **21**, 197–212 (2020).
1069 55. Bolenz, F., Kool, W., Reiter, A. M. & Eppinger, B. Metacontrol of decision-making
1070 strategies in human aging. *eLife* **8**, e49154 (2019).
1071 56. Lighthall, N. R., Pearson, J. M., Huettel, S. A. & Cabeza, R. Feedback-Based
1072 Learning in Aging: Contributions and Trajectories of Change in Striatal and
1073 Hippocampal Systems. *J. Neurosci.* **38**, 8453–8462 (2018).

1074 57. Daniel, R., Radulescu, A. & Niv, Y. Intact Reinforcement Learning But Impaired
1075 Attentional Control During Multidimensional Probabilistic Learning in Older
1076 Adults. *J. Neurosci.* **40**, 1084–1096 (2020).
1077 58. Nieuwenhuis, S. *et al.* A computational account of altered error processing in
1078 older age: dopamine and the error-related negativity. *Cogn. Affect. Behav.*
1079 *Neurosci.* **2**, 19–36 (2002).
1080 59. Karrer, T. M., Josef, A. K., Mata, R., Morris, E. D. & Samanez-Larkin, G. R.
1081 Reduced dopamine receptors and transporters but not synthesis capacity in
1082 normal aging adults: a meta-analysis. *Neurobiol. Aging* **57**, 36–46 (2017).
1083 60. Seaman, K. L. *et al.* Differential regional decline in dopamine receptor availability
1084 across adulthood: Linear and nonlinear effects of age. *Hum. Brain Mapp.* **40**,
1085 3125–3138 (2019).
1086 61. Mayr, U. & Freund, A. M. Do We Become More Prosocial as We Age, and if So,
1087 Why?: *Curr. Dir. Psychol. Sci.* (2020) doi:10.1177/0963721420910811.
1088 62. Seaman, K. L. *et al.* Adult age differences in decision making across domains:
1089 Increased discounting of social and health-related rewards. *Psychol. Aging* **31**,
1090 737–746 (2016).
1091 63. Rademacher, L., Salama, A., Gründer, G. & Spreckelmeyer, K. N. Differential
1092 patterns of nucleus accumbens activation during anticipation of monetary and
1093 social reward in young and older adults. *Soc. Cogn. Affect. Neurosci.* **9**, 825–831
1094 (2014).
1095 64. Gorlick, M. A. *et al.* Attenuating age-related learning deficits: Emotional valenced
1096 feedback interacts with task complexity. *Emotion* **13**, 250–261 (2013).

1097 65. Lockwood, P. L., O'Nell, K. C. & Apps, M. A. J. Anterior cingulate cortex: A brain
1098 system necessary for learning to reward others? *PLOS Biol.* **18**, e3000735
1099 (2020).
1100 66. Basile, B. M., Schafroth, J. L., Karaskiewicz, C. L., Chang, S. W. C. & Murray, E.
1101 A. The anterior cingulate cortex is necessary for forming prosocial preferences
1102 from vicarious reinforcement in monkeys. *PLOS Biol.* **18**, e3000677 (2020).
1103 67. Izuma, K. The social neuroscience of reputation. *Neurosci. Res.* **72**, 283–288
1104 (2012).
1105 68. Andreoni, J. Giving with impure altruism: Applications to charity and ricardian
1106 equivalence. *J. Polit. Econ.* **97**, 1447–1447 (1989).
1107 69. Andreoni, J. Impure altruism and donations to public goods: A theory of warm-
1108 glow giving. *Econ. J.* **100**, 464–477 (1990).
1109 70. Kuss, K. *et al.* A reward prediction error for charitable donations reveals outcome
1110 orientation of donators. *Soc. Cogn. Affect. Neurosci.* **8**, 216–223 (2013).
1111 71. Mobbs, D. *et al.* A key role for similarity in vicarious reward. *Science* **324**, 900–
1112 900 (2009).
1113 72. Camerer, C. F. & Fehr, E. Measuring social norms and preferences using
1114 experimental games: A guide for social scientists. *Found. Hum. Sociality Econ.*
1115 *Exp. Ethnogr. Evid. Fifteen Small-Scale Soc.* **97**, 55–95 (2004).
1116 73. Contreras-Huerta, L. S., Lockwood, P. L., Bird, G., Apps, M. A. J. & Crockett, M.
1117 J. Prosocial behavior is associated with transdiagnostic markers of affective
1118 sensitivity in multiple domains. *Emot. Wash. DC* (2020)
1119 doi:10.1037/emo0000813.

1120 74. Lockwood, P. L., Ang, Y.-S., Husain, M. & Crockett, M. J. Individual differences in
1121 empathy are associated with apathy-motivation. *Sci. Rep.* **7**, 17293 (2017).
1122 75. Venegas, J. & Clark, E. Wechsler Test of Adult Reading. in *Encyclopedia of*
1123 *Clinical Neuropsychology* (eds. Kreutzer, J. S., DeLuca, J. & Caplan, B.) 2693–
1124 2694 (Springer, 2011). doi:10.1007/978-0-387-79948-3_1500.
1125 76. Dykiert, D. & Deary, I. J. Retrospective validation of WTAR and NART scores as
1126 estimators of prior cognitive ability using the Lothian Birth Cohort 1936. *Psychol.*
1127 *Assess.* **25**, 1361–1366 (2013).
1128 77. Gordts, S., Uzieblo, K., Neumann, C., Van den Bussche, E. & Rossi, G. Validity
1129 of the Self-Report Psychopathy Scales (SRP-III Full and Short Versions) in a
1130 Community Sample. *Assessment* **24**, 308–325 (2017).
1131 78. Neumann, C. S. & Pardini, D. Factor Structure and Construct Validity of the Self-
1132 Report Psychopathy (SRP) Scale and the Youth Psychopathic Traits Inventory
1133 (YPI) in Young Men. *J. Personal. Disord.* **28**, 419–433 (2012).
1134 79. Crockett, M. J. *et al.* Dissociable Effects of Serotonin and Dopamine on the
1135 Valuation of Harm in Moral Decision Making. *Curr. Biol.* **25**, 1852–1859 (2015).
1136 80. Lockwood, P. L. *et al.* Prosocial apathy for helping others when effort is required.
1137 *Nat. Hum. Behav.* **1**, 0131 (2017).
1138 81. Lockwood, P. L. & Klein-Flügge, M. C. Computational modelling of social
1139 cognition and behaviour—a reinforcement learning primer. *Soc. Cogn. Affect.*
1140 *Neurosci.* (2020) doi:10.1093/scan/nsaa040.
1141 82. MacKay, D. J. C. *Information Theory, Inference and Learning Algorithms*.
1142 (Cambridge University Press, 2003).

83. Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J. & Friston, K. J. Bayesian Model Selection for Group Studies. *NeuroImage* **46**, 1004–1017 (2009).

84. Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P. & Dolan, R. J. Model-Based Influences on Humans' Choices and Striatal Prediction Errors. *Neuron* **69**, 1204–1215 (2011).

85. Palminteri, S., Khamassi, M., Joffily, M. & Coricelli, G. Contextual modulation of value signals in reward and punishment learning. *Nat. Commun.* **6**, 8096 (2015).

86. Palminteri, S., Wyart, V. & Koechlin, E. The Importance of Falsification in Computational Cognitive Modeling. *Trends Cogn. Sci.* **21**, 425–433 (2017).

87. R Core Team. *R: A language and environment for statistical computing. R Foundation for Statistical Computing*. (2017).

88. RStudio Team. *RStudio: Integrated Development for R. RStudio, Inc.* (2015).

89. Koller, M. robustlmm: An R Package for Robust Estimation of Linear Mixed-Effects Models. *J. Stat. Softw.* **75**, 1–24 (2016).

90. Bates, D., Mächler, M., Bolker, B. & Walker, S. Fitting Linear Mixed-Effects Models Using lme4. *J. Stat. Softw.* **67**, 1–48 (2015).

91. Kassambara, A. *rstatix: Pipe-Friendly Framework for Basic Statistical Tests. R package version 0.5.0.* (2020).

92. JASP Team. *JASP (Version 0.12.2)*. (2020).

93. Marsman, M. & Wagenmakers, E.-J. Bayesian benefits with JASP. *Eur. J. Dev. Psychol.* **14**, 545–555 (2017).

94. Tingley, D., Yamamoto, T., Hirose, K., Keele, L. & Imai, K. mediation: R Package for Causal Mediation Analysis. *J. Stat. Softw.* **59**, 1–38 (2014).

95. Venables, W. N. & Ripley, B. D. *Modern Applied Statistics with S*. (Springer, 2002).

96. Imai, K., Keele, L. & Tingley, D. A general approach to causal mediation analysis. *Psychol. Methods* **15**, 309–334 (2010).

97. Preacher, K. J., Rucker, D. D. & Hayes, A. F. Addressing Moderated Mediation Hypotheses: Theory, Methods, and Prescriptions. *Multivar. Behav. Res.* **42**, 185–227 (2007).

98. Judd, C. M. & Kenny, D. A. Process Analysis: Estimating Mediation in Treatment Evaluations. *Eval. Rev.* **5**, 602–619 (1981).