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, 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
Patricia Lockwood: P.L.Lockwood@bham.ac.uk
Centre for Human Brain Health, University of Birmingham, United Kingdom.

Abstract

Reinforcement learning is a fundamental mechanism displayed by many species from mice to humans. However, adaptive behaviour depends not only on learning associations between actions and outcomes that affect ourselves, but critically, also outcomes that affect other people. Existing studies suggest reinforcement learning ability declines across the lifespan and self-relevant learning can be computationally separated from learning about rewards for others, yet how older adults learn what rewards others is unknown. Here, using computational modelling of a probabilistic reinforcement learning task, 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 computational model with separate learning rates best explained how people learn associations for different recipients. Young adults were faster to learn when their actions benefitted themselves, compared to when they helped others. Strikingly however, older adults showed reduced self-bias, with a relative increase in the rate at which they learnt about actions that helped others, compared to themselves. Moreover, we find evidence that these group differences are associated with changes in psychopathic traits over the lifespan. In older adults, psychopathic traits were significantly reduced and negatively correlated with prosocial learning rates. Importantly, older people with the lowest levels of psychopathy had the highest prosocial learning rates. These findings suggest learning how our actions help others is preserved across the lifespan with implications for our understanding of reinforcement learning mechanisms and theoretical accounts 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[8]. It is largely unknown how ageing affects social functioning, despite the critical importance of this question. Social isolation has been found to be as damaging to physical health as smoking or excessive drinking[9]. Prosocial behaviours, or actions intended to benefit others, have a key role in maintaining vital social interactions and relationships[10] across the lifespan. In addition to the benefits for others, prosociality has been linked to improved life satisfaction[11], mental wellbeing[12], and physical health[13] 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[14]. 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[15].
Although essential for successful adaptive behaviour, some studies have found that our propensity for reinforcement learning declines in later life[8]. For example, Mell et

al. (2005) showed that older adults are impaired at learning and reversing action-outcome associations, compared to younger adults. This decline could be due to the significant age-related decrease in dopamine transmission[17–19], which has been suggested to code prediction errors[2,20,21]. Indeed, one study showed that administering L-DOPA, a dopamine precursor, to older adults increased their learning rate[22]. 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[23,24]. There is also evidence of an age-related increase in charitable donations to individuals in need[25]. At work, older adults engage in more prosocial behaviours than younger adults, according to both self-report data and colleagues' ratings[26]. Finally, self-reported altruism and decisions to donate to others have been shown to increase with age[27]. 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[28]. Nevertheless, if older adults are indeed more prosocial, we might expect that whilst self-relevant learning declines with ageing, prosocial learning could be preserved.

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[29,30]. Psychopathy has dysfunctional affective-interpersonal features at its core[31,32] but is also characterised by lifestyle and antisocial traits[33]. At the extreme, psychopathy is a severe personality condition linked to poor life outcomes, violence, and criminality[34–36]. However, several studies suggest similar behavioural and neural profiles between community samples with high levels of psychopathic traits[37] and those with clinical diagnoses of psychopathy, consistent with the Research Domains of Criteria (RDoC) conceptualisation of a dimensional approach to psychiatry[38].

Intriguingly, ageing is also associated with changes in psychopathic traits. Criminal activity increases during adolescence then declines in older adulthood[39]. Antisocial and aggressive behaviours also significantly decrease in older age, with young adults (age 16-24 years) having the highest rates of homicide[40]. Even within violent male offenders, psychopathic traits linked to an antisocial lifestyle are negatively correlated with age[41]. In community samples, ageing is associated with a decrease in both the antisocial lifestyle and affective interpersonal elements of psychopathic traits[42].

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 availability. 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). These group differences were associated with changes in psychopathic traits over the lifespan. In older adults,

psychopathic traits were significantly reduced and 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 this change is associated with a decline in psychopathic traits. These findings suggest learning how our actions help others is preserved across the lifespan.

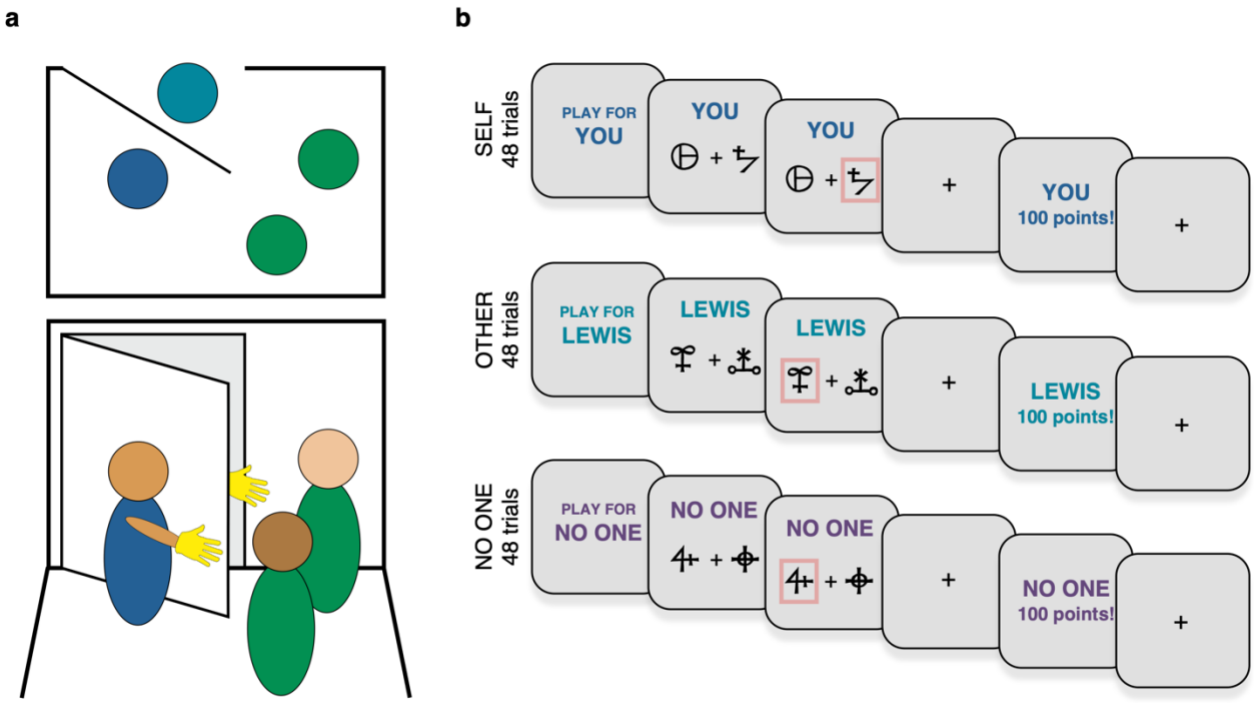

Figure 1. Prosocial learning task and social role assignment. (a) The role assignment procedure involved the participant (dark blue), confederate (light blue) and two experimenters (green). Top: from above showing the positioning of the participant and two experimenters inside the testing room, and the 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'. After this procedure they 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 trait psychopathy measure (see Methods). To ensure comparability, older adults with dementia, as diagnosed by the Addenbrooke's Cognitive Examination (ACE)[43], 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)[44]. 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.[45] and Wittmann et al.[46] (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).

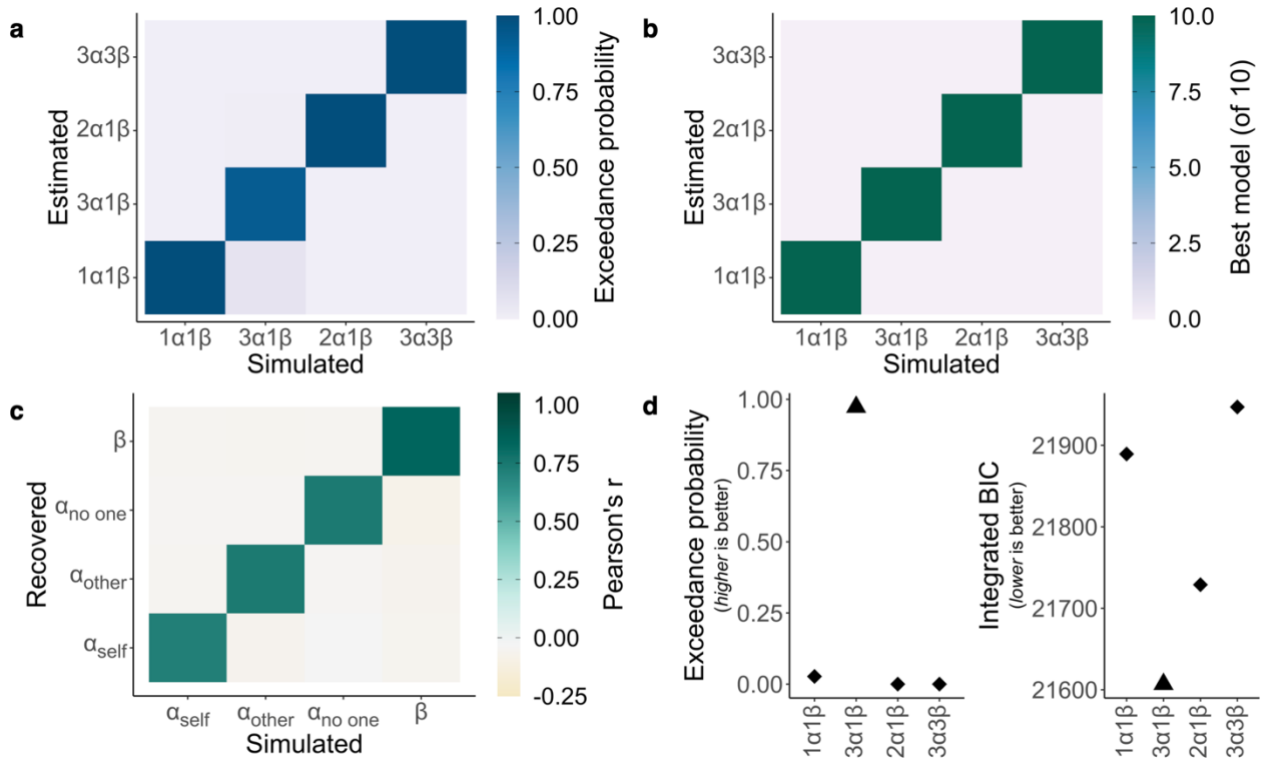

Figure 2. Model identifiability and parameter recovery. (a) Model identifiability average exceedance probability confusion matrix and **(b)** model identifiability best model selection confusion matrix. Data were simulated from 150 synthetic participants with each of our four models then Bayesian model 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.

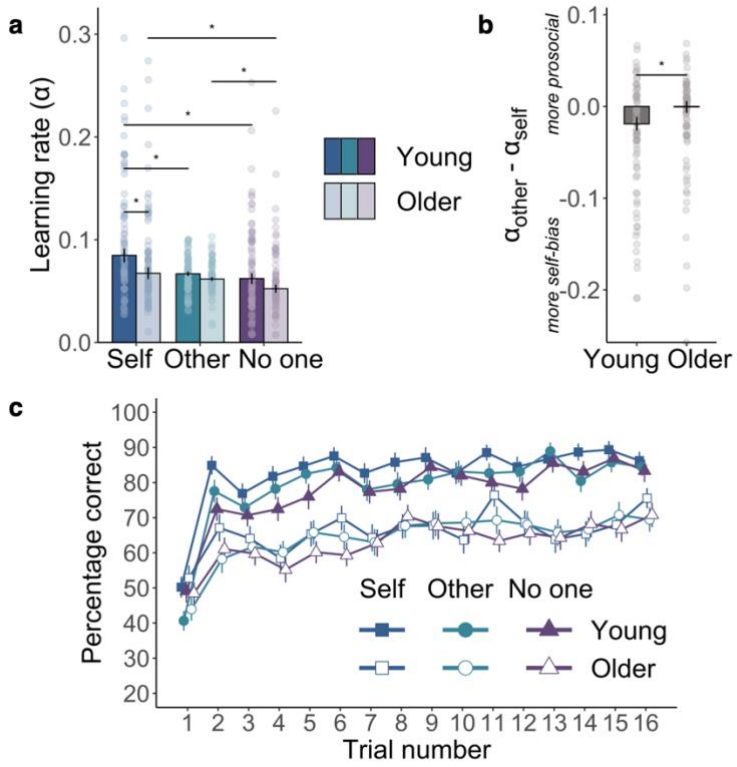

Figure 3. Age-group differences in accuracy and learning rates. (a) Comparison of learning rates from the computational model. Bars show group median. Between-group and within-group Wilcoxon t-tests show older adult's prosocial learning was preserved: learning rates in the other condition did not differ from the self condition or from young adult's prosocial learning rates. Only young participants showed self-bias, n=150 (75 young, 75 older). **(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. Asterisks represent significant differences (p<0.05). **(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).

Older adults show a reduced self-bias in learning rates

Next, we used this validated 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.

Across age groups, participants showed a higher learning rate when rewards were for themselves, compared to for another person (*b*=-0.024, 95% confidence interval [-0.034, -0.014], *z*=-4.79, *p*<0.001). Importantly however, this pattern differed between age groups. The difference between self and other learning rates was reduced in older adults, consistent with a reduced self-bias (recipient [self vs. other] * age group interaction: *b*=0.016 [0.002, 0.03], *z*=2.29, *p*=0.022). To explore this interaction, we ran planned between-group comparisons between the older and young group, in each recipient condition. Comparing α_{self} between age groups showed older adults learnt significantly more slowly for themselves (W=3512, Z=-2.6, $r_{(150)}$=0.21 [0.06, 0.73], *p*=0.009). In contrast, prosocial learning was maintained in older adults, with a Bayes factor suggesting strong evidence in favour of no difference in α_{other}, compared to young adults (W=3042, Z=-0.86, $r_{(150)}$=0.07 [0.003, 0.23], *p*=0.39, BF_{01}=4.26). As expected, across age groups learning was slower for no one than self (recipient [self vs. no one] main effect: *b*=-0.023 [-0.033, -0.013], *z*=-4.57, *p*<0.001). $\alpha_{no\ one}$ did not significantly differ between groups, although there was not strong evidence in support of the null (W=3241, Z=-1.6, $r_{(150)}$=0.13 [0.008, 0.29], *p*=0.11, BF_{01}=2.04).

We also ran planned within-subject comparisons between recipients for each age group separately. As predicted, young adults learnt preferentially for themselves (vs. α_{other} V=659, Z=-4.0, $r_{(75)}$=0.47 [0.27, 0.63], *p*<0.001; vs. $\alpha_{no\ one}$ V=928, Z=-2.62, $r_{(75)}$=0.30 [0.09, 0.51], *p*=0.009) and did not differentiate between another person and no one (V=1533, Z=-0.57, $r_{(75)}$=0.067 [0.003, 0.30], *p*=0.57, BF_{01}=5.08). In contrast, for older adults, there was no significant difference between α_{self} and α_{other} (V=1150, Z=-1.4, $r_{(75)}$=0.17 [0.01, 0.37], *p*=0.15, BF_{01}=1.08) but learning rates were slower for no one (vs. α_{self} V=901, Z=-2.8, $r_{(75)}$=0.32 [0.10, 0.52], *p*=0.006; vs. α_{other} V=976, Z=-

2.4, $r_{(75)}$=0.27 [0.05, 0.50], *p*=0.02). This demonstrates specificity of the age-related change to prosocial learning. In other words, older adults showed a relative increase in the willingness to learn about actions that helped others. Finally, we also observed a main effect of age on both learning rates and temperature parameters. Older adults showed slower learning overall compared to younger adults (*b*=-0.019 95% CI [-0.028, -0.009], Z=-3.73, *p*<0.001) and higher levels of exploration of choice options (median β young: 0.05, older: 0.19, W=1511, Z=-4.9, $r_{(150)}$=0.4 [0.25, 0.54], *p*<0.001).

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 no 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 (*b*s<0.16, *z*s<0.96, *p*s>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[37]. Critically, psychopathic traits are most closely linked to alterations in social behaviour and willingness to help others. Therefore, we also asked participants to complete the Self-report Psychopathy Scale[33], a measure of psychopathic traits that robustly capture its latent structure (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 affective-interpersonal (W=3558, Z=-3.1, $r_{(148)}$=0.26 [0.11, 0.41], *p*=0.002) and the lifestyle-antisocial subscales (W=3471, Z=-2.8, $r_{(148)}$=0.23 [0.07, 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). Intriguingly, this relationship was significantly more negative (Z=3.3, *p*=0.001) than the equivalent correlation in young adults, which had a positive sign ($r_{s(74)}$=0.21 [-0.02, 0.42], *p*=0.07; Figure 4a). This pattern of results was the same when correlating the relative difference between α_{other} - α_{self} (rather than the raw α_{other}) with core psychopathic traits (Figure 4b). 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 3). The negative relationship between α_{other} and core psychopathic traits for older adults also remained significant after applying false discovery rate correction for multiple comparisons across this correlation and the five other age group & recipient combinations (see Supplementary Table 4). Moreover, the finding that no significant correlations were apparent between psychopathy and α_{self} or $\alpha_{no\ one}$ (*p*s>0.15; Supplementary Table 4) suggests a specificity in the relevance of psychopathy to prosocial learning in older adults.

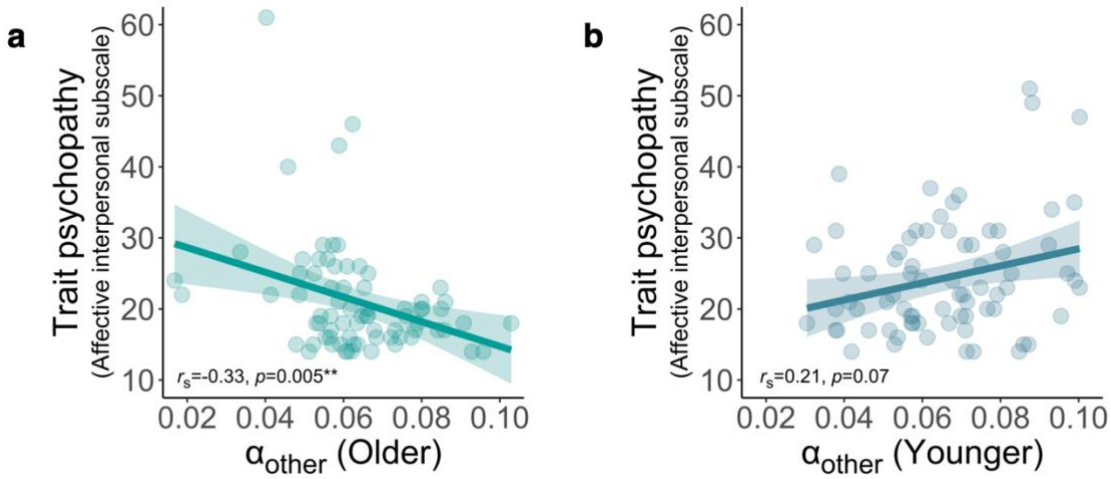

Figure 4. Correlations between prosocial learning rates (α_{other}) and scores on the affective interpersonal subscale of the Self Report Psychopathy scale. **(a)** For older adults, trait psychopathy levels are negatively correlated with prosocial learning rates (r_s=-0.33 [-0.52, -0.11], p=0.005). **(b)** There is no significant relationship for young adults (r_s=0.21 [-0.02, 0.42], p=0.07) and the correlation is significantly more negative (Z=3.3, p=0.001) for older than young adults. This pattern of results is the same when considering correlations between trait psychopathy and the lack of self-bias in learning (α_{other} - α_{self}; not shown). This measure of prosocial learning is also negatively correlated with trait psychopathy scores 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=2.1, p=0.03).

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[47] and our actions affect outcomes for others. Here, for the first time, we apply 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 a relative increase in their

prosocial learning rates, ageing was also associated with lower psychopathic traits, which were specifically linked to prosocial learning ability.

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 or no one received the reward. However, this self-bias was reduced in older adults, who showed a relative increase in prosocial learning rates. 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, learning rates were not associated with executive function or an age-standardised measure of intelligence. We also show that our results remain the same after controlling for these measures. These findings suggest that the observed decline in self-relevant learning rates, but relative increase in prosocial learning rates, for older adults are not explained by changes in these broad abilities. Considering learning more specifically, a recent comparison of motivation and ability as explanations for ageing-related reductions in model-based strategies during self-relevant learning supported the limited cognitive abilities account[48]. Our finding that learning rates for self-relevant outcomes were reduced in older adults is in line with a degeneration in the neurocognitive systems required for successful learning. Research combining models of learning with neuroimaging and pharmacological manipulations suggests ageing reduces the ability to generate reward prediction errors[49] due to declines in dopamine functioning[22,50]. Differences in motivation could also be applicable for self-

relevant learning as the subjective value of financial outcomes is also likely to decrease in older age, due to changes in wealth across the lifespan[28].

Our findings suggest that despite declines in learning ability associated with ageing, motivation could play a role in preserving learning to help others. This interpretation aligns with an emerging literature showing older adults may be more prosocial and less self-biased than younger adults[23,24,51]. The assertion that relatively preserved prosocial learning is related to increased prosocial motivation is supported by our observed link between learning rates and psychopathic traits. Psychopathic traits were significantly reduced in our older adult sample, dovetailing with similar previous findings on this trait[41,42] and broader trait benevolence[27]. Importantly, we found psychopathic traits in older adults negatively correlated with prosocial learning rates. Self-bias in learning rates (i.e. higher learning rate for self compared to other) was most reduced, and even reversed, in the older people lowest on psychopathic traits. Notably, this negative correlation between psychopathy and prosocial learning was only found for older adults. This suggests that age-related differences in prosocial learning could be linked to basic shifts in individual traits and motivations over the lifespan, not just to domain-general reductions in cognitive abilities.

Moreover, our task included a control condition where no one benefitted. This was important 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 (if anything they learned faster for 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 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[7,52], and causally in monkeys with focal lesions[53].

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[54], the 'warm glow' of helping[55–57], or vicarious reinforcement from positive outcomes for others[58]. 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 relative increase in prosocial 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[59] 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).

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. Further research with participants of all ages is required to address these outstanding questions. Moreover, 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 through university databases, social media and the community. This sample size gave us 88% power to detect a medium size effect (d=0.5). Exclusion criteria included 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)[43].

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 approved by the Oxford University Medical Sciences Inter Divisional Research Ethics Committee and National Health Service Ethics.

Prosocial learning task

The prosocial learning task is a probabilistic reinforcement learning task, with rewards in one of three recipient conditions: for the participant themselves (*self*), for another participant (*other*; prosocial condition), and for *no one* (control condition). Each trial presents two symbols, one associated with a high (75%) probability of gaining points and the other with a low (25%) probability of gaining points. The two symbols were randomly assigned to the left or right side of the screen and selected via a corresponding button press. Participants select a symbol then receive feedback on whether they obtained points or not (see Figure 1b) so learn over time which symbol maximises rewards. The experiment was subdivided into blocks, i.e. 16 trials pairings the same two symbols for the same recipient. Participants completed three blocks, a total of 48 trials, in each recipient condition, resulting in 144 trials overall (see Supplementary Information for trial structure). Blocks for different recipients were pseudo-randomly ordered such that the same recipient block did not occur twice in a row.

On trials in the *self* condition, points translated into increased payment for the participant themselves. These blocks started with “play for you” displayed and had the word “you” at the top of each screen. Blocks in the *no one* condition had “no one” in place of “you” and points were not converted into any extra payment for anyone. In the prosocial ‘*other’* condition, participants earned points that translated into additional payment for a second participant, actually a confederate. Participants were told that this payment would be given anonymously, they would never meet the other person, and that the person was not even aware of them completing this task (see Supplementary Information). The name of the confederate, gender-matched to the participant, was displayed on these blocks at the start and on each screen (Figure 1b). Thus, participants were explicitly aware who their decisions affected on each trial. Stimuli were presented using Presentation (Neurobehavioral Systems – https://www.neurobs.com/).

Questionnaire measures

Dementia screening and executive function. Older adults were screened for dementia using the Addenbrooke's Cognitive Examination (ACE-III)[43]. The ACE examines five cognitive domains; attention, memory, language, fluency and visuospatial abilities. The ACE-III is scored out of 100 and as a screening tool, a cut-off score of 82/100 denotes significant cognitive impairment. We also used scores on the attention and memory domains in control analyses as proxies for executive function in older adults.

General intelligence. All participants completed the Wechsler Test of Adult Reading (WTAR)[44] as a measure of IQ. The WTAR requires participants to pronounce 50 words that have unusual grapheme-to-phoneme translation. This means the test measures reading recognition or existing knowledge of the words, rather than ability to apply rules for pronunciation. The WTAR was developed and standardised with the Wechsler Memory and Adult Intelligence Scales and correlates highly with these measures[60]. Standardisation involves adjusting for healthy age-related differences. The test is suitable for participants aged 16-89, covering our full sample, and scores in older age have been shown to correlate with cognitive ability earlier in life[61].

Psychopathic traits. Participants completed the short form of the Self-Report Psychopathy Scale 4th Edition (SRP-IV-SF)[33]. This scale consists of 29 items, 7 each measuring: interpersonal, affective, lifestyle and antisocial tendencies (plus 'I have been convicted of a serious crime'). We used the two-factor grouping, summing the core, affective-interpersonal items and separately, the lifestyle-antisocial items for use in analysis. The robust psychometric properties of this measure have been established in community[62] and offender populations through construct and convergent validity[63], internal consistency, and reliability[33].

Procedure

Role assignment. To enhance belief that points earned in the prosocial recipient benefitted another person, we conducted a role assignment procedure based on set-up used in several studies of social decision-making[64,65]. Participants were instructed not to speak and wore a glove to hide their identity. A second experimenter brought the confederate, also wearing a glove, to the other side of the door. Participants only ever saw the gloved hand of the confederate, but they waved to each other to make it clear there was another person there (Figure 1a). The experimenter tossed a coin to determine who picked a ball from the box first and then told the participants which roles they had been assigned to, based on the ball they picked. Our procedure ensured that participants always ended up in the role of the person performing the prosocial learning experiment. Participants were unaware of the age of the other person, but the experimenter used a name for them suggesting their gender was the same as the participant.

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

Computational modelling

We modelled learning during the task with a reinforcement learning algorithm[14], creating variations of the models through the number of parameters used to explain the learning rate and temperature parameters in the task[66]. The basis of the reinforcement learning algorithm is the expectation that an action (or stimulus) a will provide reward on the following trial. This expected value, $Q_{t+1}(a)$ is quantified as a

function of current expectations $Q_t(a)$ and the prediction error δ_t, which is scaled by the learning rate α:

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

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

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

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

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

Model fitting

We used MATLAB 2019b (The MathWorks Inc) for all model fitting and comparison. To fit the variations of the learning model (see below) to (real and simulated) participant data we used an iterative maximum a posteriori (MAP) approach previously described[45,46]. This method provides a better estimation than a single-step maximum likelihood estimation (MLE) alone by being less susceptible to the influence of outliers. It does this via implementing two levels: the lower level of the individual subjects and the higher-level reflecting our full sample. For the real participant data, we fit the model across groups to provide the most conservative comparison, so this full sample combined young and older participants.

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

Model comparison

Our hypotheses generated four models to compare which differed in whether the model parameters (α and β) for each participant had one value across recipient conditions or depended on the recipient (self, other and no one). For model comparison, we calculated the Laplace approximation of the log model evidence (more positive values indicating better model fit; MacKay, 2003) and submitted these to a random-effects analysis using the spm_BMS routine[68] from SPM 8 (http://www.fil.ion.ucl.ac.uk/spm/software/spm8/). This generates the exceedance probability: the posterior probability that each model is the most likely of the model set in the population (higher is better, over .95 indicates strong evidence in favour of a model). For the models of real participant data, we also calculated the integrated BIC (lower is better[45,46]) and R^2 as additional measures of model fit. To calculate the model R^2, we extracted the choice probabilities generated for each participant on each trial from the winning model. We then took the squared median choice probability across participants. The $3\alpha1\beta$ model had the best evidence on all measures (see Supplementary Table 2).

Simulation experiments

We used simulation experiments to assess that our experiment allowed us to dissociate models of interest, as well as parameters of interest within the winning model. We simulated data from all four models to establish that our model comparison procedure (see above) could accurately identify the best model across a wide range of parameter values. For this model identifiability analysis, we simulated data from 150 participants, drawing parameters from distributions commonly used in the reinforcement learning literature[69,70]. Learning rates (α) were drawn from a beta distribution (betapdf(parameter,1.1,1.1)) and softmax temperature parameters (β) from a gamma distribution (gampdf(parameter,1.2,5)). We fitted the models to this simulated data set using the same MAP process as applied to the experimental participants' data and repeated this whole procedure 10 times. By plotting the confusion matrices of average exceedance probability (across the 10 runs; Figure 2a) and how many times each model won (Figure 2b), we show the models are identifiable using our model comparison process.

Our winning model contained 4 free parameters (α_{self}, α_{other}, $\alpha_{no\ one}$, β). To assess the 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[66,71]. 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.

Statistical analysis

Analysis of group and recipient differences in the fitted model parameters and behavioural data was run in R[72] with R Studio[73]. Correlations of learning rates with trait psychopathy (Supplementary Table 3 & 4) and neuropsychological measures (Supplementary Table 5) were calculated with Spearman's Rho nonparametric tests. We used robust linear mixed-effect models (*rlmer* function; robustlmm package[74] to

predict learning rates and generalised linear mixed-effects models (*glmer* function; lme4 package[75]) for the trial-by-trial data (binary outcome of choosing the high vs. low reward option). Each included 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 6). To control for IQ and executive function in the associations between older adults' prosocial learning rates and trait psychopathy, we ran partial correlations each controlling for one of WTAR, ACE memory and ACE attention scores (Supplementary Table 7).

For simple and post hoc comparisons, we used nonparametric tests when measures violated normality assumptions, which included learning rates. 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[76]. Bayes factors (BF_{01}) for non-significant results were calculated using nonparametric paired and independent t-tests in JASP[77] 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[78].

Data availability

Data are available at: https://osf.io/xgw7h/.

Code availability

Code for modelling and analysis is available at: https://osf.io/xgw7h/.

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. Samanez-Larkin, G. R. & Knutson, B. Decision making in the ageing brain: changes in affective and motivational circuits. *Nat. Rev. Neurosci.* **16**, 278–289 (2015).
9. Holt-Lunstad, J., Smith, T. B., Baker, M., Harris, T. & Stephenson, D. Loneliness and social isolation as risk factors for mortality: a meta-analytic review. *Perspect. Psychol. Sci. J. Assoc. Psychol. Sci.* **10**, 227–237 (2015).
10. Fehr, E. & Fischbacher, U. The nature of human altruism. *Nature* **425**, 785–791 (2003).
11. Buchanan, K. E. & Bardi, A. Acts of Kindness and Acts of Novelty Affect Life Satisfaction. *J. Soc. Psychol.* **150**, 235–237 (2010).
12. Raposa, E. B., Laws, H. B. & Ansell, E. B. Prosocial Behavior Mitigates the Negative Effects of Stress in Everyday Life. *Clin. Psychol. Sci.* **4**, 691–698 (2016).
13. Post, S. G. Altruism, Happiness, and Health: It's Good to Be Good. *Int. J. Behav. Med.* **12**, 66–77 (2005).

14. Sutton, R. S. & Barto, A. G. *Reinforcement learning: An introduction*. (MIT Press, 2011).
15. Schultz, W. Multiple dopamine functions at different time courses. *Annu. Rev. Neurosci.* **30**, 259–288 (2007).
16. Mell, T. *et al.* Effect of aging on stimulus-reward association learning. *Neuropsychologia* **43**, 554–563 (2005).
17. Bäckman, L., Nyberg, L., Lindenberger, U., Li, S.-C. & Farde, L. The correlative triad among aging, dopamine, and cognition: Current status and future prospects. *Neurosci. Biobehav. Rev.* **30**, 791–807 (2006).
18. Dreher, J.-C., Meyer-Lindenberg, A., Kohn, P. & Berman, K. F. Age-related changes in midbrain dopaminergic regulation of the human reward system. *Proc. Natl. Acad. Sci.* **105**, 15106–15111 (2008).
19. Li, S.-C., Lindenberger, U. & Bäckman, L. Dopaminergic modulation of cognition across the life span. *Neurosci. Biobehav. Rev.* **34**, 625–630 (2010).
20. Schultz, W. Updating dopamine reward signals. *Curr. Opin. Neurobiol.* **23**, 229–238 (2013).
21. Schultz, W. Dopamine reward prediction-error signalling: a two-component response. *Nat. Rev. Neurosci.* **17**, 183–195 (2016).
22. Chowdhury, R. *et al.* Dopamine restores reward prediction errors in old age. *Nat. Neurosci.* **16**, 648–653 (2013).
23. Engel, C. Dictator games: a meta study. *Exp. Econ.* **14**, 583–610 (2011).
24. Matsumoto, Y., Yamagishi, T., Li, Y. & Kiyonari, T. Prosocial Behavior Increases with Age across Five Economic Games. *PLOS ONE* **11**, e0158671 (2016).
25. Sze, J. A., Gyurak, A., Goodkind, M. S. & Levenson, R. W. Greater Emotional Empathy and Prosocial Behavior in Late Life. *Emotion* **12**, 1129–1140 (2012).
26. Ng, T. W. H. & Feldman, D. C. The relationship of age to ten dimensions of job performance. *J. Appl. Psychol.* **93**, 392–423 (2008).
27. Hubbard, J., Harbaugh, W. T., Srivastava, S., Degras, D. & Mayr, U. A general benevolence dimension that links neural, psychological, economic, and life-span data on altruistic tendencies. *J. Exp. Psychol. Gen.* **145**, 1351–1358 (2016).
28. Cheung, F. & Lucas, R. E. When Does Money Matter Most? Examining the Association between Income and Life Satisfaction over the Life Course. *Psychol. Aging* **30**, 120–135 (2015).

29. Buckholtz, J. W. *et al.* Mesolimbic Dopamine Reward System Hypersensitivity in Individuals with Psychopathic Traits. *Nat. Neurosci.* **13**, 419–421 (2010).
30. White, B. A. Who cares when nobody is watching? Psychopathic traits and empathy in prosocial behaviors. *Personal. Individ. Differ.* **56**, 116–121 (2014).
31. White, S. F. *et al.* Reduced Amygdala Response in Youths With Disruptive Behavior Disorders and Psychopathic Traits: Decreased Emotional Response Versus Increased Top-Down Attention to Nonemotional Features. *Am. J. Psychiatry* **169**, 750–758 (2012).
32. Marsh, A. A. *et al.* Empathic responsiveness in amygdala and anterior cingulate cortex in youths with psychopathic traits. *J. Child Psychol. Psychiatry* **54**, 900–910 (2013).
33. Paulhus, D. L., Neumann, C. S., Hare, R. D., Williams, K. M. & Hemphill, J. F. *Self-report Psychopathy Scale 4th Edition (SRP 4) Manual*. (MHS, Multi-Health Systems Incorporated, 2017).
34. Leistico, A.-M. R., Salekin, R. T., DeCoster, J. & Rogers, R. A large-scale meta-analysis relating the Hare measures of psychopathy to antisocial conduct. *Law Hum. Behav.* **32**, 28–45 (2008).
35. Blais, J., Solodukhin, E. & Forth, A. E. A Meta-Analysis Exploring the Relationship Between Psychopathy and Instrumental Versus Reactive Violence. *Crim. Justice Behav.* **41**, 797–821 (2014).
36. Asscher, J. J. *et al.* The relationship between juvenile psychopathic traits, delinquency and (violent) recidivism: A meta-analysis. *J. Child Psychol. Psychiatry* **52**, 1134–1143 (2011).
37. Seara-Cardoso, A., Neumann, C., Roiser, J., McCrory, E. & Viding, E. Investigating associations between empathy, morality and psychopathic personality traits in the general population. *Personal. Individ. Differ.* **52**, 67–71 (2012).
38. Insel, T. R. The NIMH Research Domain Criteria (RDoC) Project: Precision Medicine for Psychiatry. *Am. J. Psychiatry* **171**, 395–397 (2014).
39. Liberman, A. *The long view of crime: A synthesis of longitudinal research*. (Springer, 2008).
40. Homicide in England and Wales - Office for National Statistics. https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/articles/

homicideinenglandandwales/yearendingmarch2018#what-do-we-know-about-suspects.

41. Huchzermeier, C. *et al.* Are there age-related effects in antisocial personality disorders and psychopathy? *J. Forensic Leg. Med.* **15**, 213–218 (2008).
42. Gill, D. J. & Crino, R. D. The Relationship between Psychopathy and Age in a Non-Clinical Community Convenience Sample. *Psychiatry Psychol. Law* **19**, 547–557 (2012).
43. Hsieh, S., Schubert, S., Hoon, C., Mioshi, E. & Hodges, J. R. Validation of the Addenbrooke's Cognitive Examination III in frontotemporal dementia and Alzheimer's disease. *Dement. Geriatr. Cogn. Disord.* **36**, 242–250 (2013).
44. The Psychological Corporation. *Wechsler test of adult reading*. (Harcourt Assessment, 2001).
45. Huys, Q. J. *et al.* Disentangling the roles of approach, activation and valence in instrumental and pavlovian responding. *PLoS Comput Biol* **7**, e1002028 (2011).
46. Wittmann, M. K. *et al.* Global reward state affects learning and activity in raphe nucleus and anterior insula in monkeys. *Nat. Commun.* **11**, 3771 (2020).
47. Olsson, A., Knapska, E. & Lindström, B. The neural and computational systems of social learning. *Nat. Rev. Neurosci.* **21**, 197–212 (2020).
48. Bolenz, F., Kool, W., Reiter, A. M. & Eppinger, B. Metacontrol of decision-making strategies in human aging. *eLife* **8**, e49154 (2019).
49. Eppinger, B., Hämmerer, D. & Li, S.-C. Neuromodulation of reward-based learning and decision making in human aging. *Ann. N. Y. Acad. Sci.* **1235**, 1–17 (2011).
50. Nieuwenhuis, S. *et al.* A computational account of altered error processing in older age: dopamine and the error-related negativity. *Cogn. Affect. Behav. Neurosci.* **2**, 19–36 (2002).
51. Mayr, U. & Freund, A. M. Do We Become More Prosocial as We Age, and if So, Why?: *Curr. Dir. Psychol. Sci.* (2020) doi:10.1177/0963721420910811.
52. Lockwood, P. L., O'Nell, K. C. & Apps, M. A. J. Anterior cingulate cortex: A brain system necessary for learning to reward others? *PLOS Biol.* **18**, e3000735 (2020).
53. Basile, B. M., Schafroth, J. L., Karaskiewicz, C. L., Chang, S. W. C. & Murray, E. A. The anterior cingulate cortex is necessary for forming prosocial preferences from vicarious reinforcement in monkeys. *PLOS Biol.* **18**, e3000677 (2020).

54. Izuma, K. The social neuroscience of reputation. *Neurosci. Res.* **72**, 283–288 (2012).
55. Andreoni, J. Giving with impure altruism: Applications to charity and ricardian equivalence. *J. Polit. Econ.* **97**, 1447–1447 (1989).
56. Andreoni, J. Impure altruism and donations to public goods: A theory of warm-glow giving. *Econ. J.* **100**, 464–477 (1990).
57. Kuss, K. *et al.* A reward prediction error for charitable donations reveals outcome orientation of donators. *Soc. Cogn. Affect. Neurosci.* **8**, 216–223 (2013).
58. Mobbs, D. *et al.* A key role for similarity in vicarious reward. *Science* **324**, 900–900 (2009).
59. Camerer, C. F. & Fehr, E. Measuring social norms and preferences using experimental games: A guide for social scientists. *Found. Hum. Sociality Econ. Exp. Ethnogr. Evid. Fifteen Small-Scale Soc.* **97**, 55–95 (2004).
60. Venegas, J. & Clark, E. Wechsler Test of Adult Reading. in *Encyclopedia of Clinical Neuropsychology* (eds. Kreutzer, J. S., DeLuca, J. & Caplan, B.) 2693–2694 (Springer, 2011). doi:10.1007/978-0-387-79948-3_1500.
61. Dykiert, D. & Deary, I. J. Retrospective validation of WTAR and NART scores as estimators of prior cognitive ability using the Lothian Birth Cohort 1936. *Psychol. Assess.* **25**, 1361–1366 (2013).
62. Gordts, S., Uzieblo, K., Neumann, C., Van den Bussche, E. & Rossi, G. Validity of the Self-Report Psychopathy Scales (SRP-III Full and Short Versions) in a Community Sample. *Assessment* **24**, 308–325 (2017).
63. Neumann, C. S. & Pardini, D. Factor Structure and Construct Validity of the Self-Report Psychopathy (SRP) Scale and the Youth Psychopathic Traits Inventory (YPI) in Young Men. *J. Personal. Disord.* **28**, 419–433 (2012).
64. Crockett, M. J. *et al.* Dissociable Effects of Serotonin and Dopamine on the Valuation of Harm in Moral Decision Making. *Curr. Biol.* **25**, 1852–1859 (2015).
65. Lockwood, P. L. *et al.* Prosocial apathy for helping others when effort is required. *Nat. Hum. Behav.* **1**, 0131 (2017).
66. Lockwood, P. L. & Klein-Flügge, M. C. Computational modelling of social cognition and behaviour—a reinforcement learning primer. *Soc. Cogn. Affect. Neurosci.* (2020) doi:10.1093/scan/nsaa040.
67. MacKay, D. J. C. *Information Theory, Inference and Learning Algorithms*. (Cambridge University Press, 2003).

68. 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).
69. 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).
70. Palminteri, S., Khamassi, M., Joffily, M. & Coricelli, G. Contextual modulation of value signals in reward and punishment learning. *Nat. Commun.* **6**, 8096 (2015).
71. Palminteri, S., Wyart, V. & Koechlin, E. The Importance of Falsification in Computational Cognitive Modeling. *Trends Cogn. Sci.* **21**, 425–433 (2017).
72. R Core Team. *R: A language and environment for statistical computing. R Foundation for Statistical Computing*. (2017).
73. RStudio Team. *RStudio: Integrated Development for R. RStudio, Inc.* (2015).
74. Koller, M. robustlmm: An R Package for Robust Estimation of Linear Mixed-Effects Models. *J. Stat. Softw.* **75**, 1–24 (2016).
75. Bates, D., Mächler, M., Bolker, B. & Walker, S. Fitting Linear Mixed-Effects Models Using lme4. *J. Stat. Softw.* **67**, 1–48 (2015).
76. Kassambara, A. *rstatix: Pipe-Friendly Framework for Basic Statistical Tests. R package version 0.5.0.* (2020).
77. JASP Team. *JASP (Version 0.12.2)*. (2020).
78. Marsman, M. & Wagenmakers, E.-J. Bayesian benefits with JASP. *Eur. J. Dev. Psychol.* **14**, 545–555 (2017).