A fixation dependent decision model of charitable choice

Philip Pärnamets

Abstract: Building on previous work showing that eye gaze plays a role during moral decision-making and that the underlying mechanisms might be characterised as a fixation dependent drift-diffusion process, donation decisions between charitable organisations were studied. Models were fit with full, no or partial fixation dependence. Results indicate the model with partial fixation dependence provided the best fit to the empirical data and could capture many aspects of the underlying choice and gaze data.

Keywords: Morality; decision-making; eye tracking; computational modelling

Understanding how decisions are reached in the moment of choice is an important goal for explaining how human decision-making works. Recent work has suggested that value-based economic decisions, such as choices between foodstuffs or consumer goods, can be modelled as a comparison process biased by eye gaze (Krajbich, Armel & Rangel, 2010; Krajbich & Rangel, 2011; Krajbich, Lu, Camerer & Rangel, 2012). This work builds on earlier findings indicating that visual attention has a causal influence on the outcome of simple choices (Shimojo, Simion, Shimojo & Scheier, 2003; Armel, Beaumel & Rangel, 2008; Milosavljevic, Navalpakkam, Koch & Rangel, 2012). There is thus an emerging understanding that attention, captured by eye gaze, tracks and influences decision-making during simple binary and trinary choices.

In the computational model proposed by Krajbich et al. (2010), the attentional drift-diffusion model (aDDM), choice is modelled as evidence accumulation proportional to the relative value of the options under consideration. In this model, the direction of gaze affects the speed of the accumulation by discounting the value of the non-fixated option. The aDDM for binary choices can be characterised by the following equation:

$$V_t = V_{t-1} + d\left(r_{fix} - \theta r_{nonfix}\right) + N(0,\sigma)$$

Where V is the decision value, which is accumulates towards either 1 or -1 and which is assumed to be starting at 0. The value of the options under consideration is given by r, indexed by fixation direction. The speed of evidence accumulation is controlled by the drift parameter d (in units ms^{-1}). The non-fixated item is discounted by θ, which can take values in the interval [0,1]. Finally, N is white Gaussian noise with variance σ^2.

Gaze and moral choice

In Pärnamets et al. (in press) it was shown that by taking eye gaze as an index of the developing decision process, it was possible to influence participants' choices in response to abstract moral questions, such as 'Is murder justifiable'. In that study, participants were asked to listen to moral questions and then choose, when prompted, the alternative, of two presented alternatives, which they felt was morally right in relation to the question. In the case of 'Is murder justifiable' alternatives could be 'Sometimes justifiable' and 'Never justifiable'. Unbeknownst to the participants, the timing of the decision prompt was dependent on their gaze patterns. The decision prompt was set to trigger when they had viewed each option a predetermined amount of time. These results showed for the first time that where participants were looking during moral choices could influence the outcome of those choices. In other work, eye gaze has also been shown to differentially support responses to complex moral dilemmas (Pärnamets, Hall & Johansson, 2014).

Building on these findings, linking eye gaze to moral choice, Pärnamets, Balkenius & Richardson (2014) investigated if the aDDM could be used to model data from moral choices similar to the ones studied in Pärnamets et al. (in press). The results of Pärnamets, Balkenius & Richardson (2014) were promising in this regard, indicating that a fixation dependent model performed better than an alternative model disregarding eye gaze (ie. a regular drift diffusion model). This suggested a practical route towards a unifying mechanistic account of evidence accumulation in the moment of choice for a wide range of separate choice domains. However, the models studied performed poorly with respect to many other aspects of the gaze data. Additionally, the relationship between value and choice probability could not be studied due to the post-hoc value sampling method used. It is therefore important to extend and replicate the previous findings to improve our understanding of how moral choices might share mechanisms with non-moral choices.

Aim

The aim of the present paper is to extend the study of the applicability of the aDDM in the moral domain by using a different stimulus set representing alternative moral choices compared to those studied previously. In the present work, instead of investigating choices between alternatives to abstract moral principles, choices between charitable organisations were studied. There are several good reasons for this. One is practical, charitable organisations are fairly concrete options which allow for prior valuation by participants before choice, which will improve the reliability of the modelling. Second, charitable choices have been frequently used to study neural mechanisms of value-based choice (e.g. Moll et al., 2006; Hare, Camerer, Knoepfle, O'Doherty & Rangel, 2010). Third, using charitable choices as stimuli entails moving away from studying hypothetical decisions and, further, that the choices which participants make can be realised, thereby probably improving the ecological validity of the task.

Methods

Empirical data

Equipment and material

Eye tracking was performed using an SMI HiSpeed eye tracker recording monocularly at 500 Hz. Stimuli were presented on a 19" screen running 1280*1024 pixels resolution using PsychoPhysics Toolbox (Kleiner, Brainard & Pelli, 2007) running on MATLAB 2012b (The MathWorks, Natick, MA.). Calibration was performed using a 13 point calibration routine followed by 4 validation points. Calibrations with error exceeding 0.75° visual angle in more than one point were rerun, resulting in an average calibration error less than 0.5°.

Stimulus material consisted of a list of 31 different charities which operated in Sweden at the time of recording. Charities were selected based on familiarity from a larger list of organisations having so called “90 numbers”, which is a donation number that only fiscally responsible charities are allowed to have.

Participants

A total of 26 participants were recruited from the student population at Lund University using ads posted on library notice boards. Of the recruited participants 10 were male and 16 were female. Average age was M = 25.1, SD = 5.5

Procedure

Participants were first asked to rate all the charitable organisations. For each organisation participants were asked how familiar they were with that organisation, how likely they would donate to that organisation and how valuable they considered that organisation’s work to be. The last measure was designated to be the value measure used to fit the models, while the first two were collected for exploratory purposes and to ensure that participants wouldn’t associate each charity with a single number during later choice.

Following a twenty minute filler task part of another experiment, the eye-tracker was calibrated to the participants. Participants were instructed that they were going to make a number of binary choices between the different charities they had previously rated. They were also told that after they had completed the choices, one of their chosen charities would be randomly selected by the computer, and displayed to them. This choice would later be realised by the experimenters in the form of a 150 SEK donation to that charity.

During the experiment, each trial was preceded with a 0.5s fixation cross, after which the names of two charities, drawn randomly from the full set, were displayed. One charity was displayed at the right hand side of the screen and the other on the left hand side. Participants could view the option as long as they wished, and indicated their choice by button press. Each participant completed 100 trials. Due to a programming error, the first five participants only completed 80 trials.

Once the experiment was completed, participants were debriefed, asked to sign consent and data release forms and paid with a cinema voucher for their participation.

Model Fitting

The models

Three models were fit to the empirical data using only the odd trials from each participant. The first model, the Free Bias model, was an aDDM model where each parameter was allowed to vary. The other two models were special cases of the first, where the θ parameter was held fixed. For the No Bias model, θ was fixed to the value of 1. This means that gaze direction has no impact on evidence accumulation in the model. In the final model, the Full Bias model, θ was fixed to 0, meaning that only the value of the currently fixated option was accumulated during each time step.

Fixations

All fixations from the odd trials of the empirical data were extracted and binned by value difference. Fixation durations within each bin were fit to a log-normal distribution. Parameters resulting from these distributions were later used to generate fixation durations during model fitting and simulation. Transitions were modelled using the empirical transition probabilities as a function of the number of consecutive fixations to the same option.

The first fixation direction was modelled as a binomial draw using the empirical probability of a leftward fixation (68.6%).

Parameter selection

The models were fit using a grid search in parameter space, first with a broad set of parameters and then followed by a narrow set. As in previous work σ was varied as function of d, letting $\sigma = d*\eta$.

For the Free Bias model, during the first search $\theta = \{0.1, 0.3, 0.5, 0.7, 0.9\}$ while $d = \{0.00005, 0.0001, 0.0002\}$ and $\eta = \{90, 110, 130, 150\}$. In the second search $\theta = \{0.3, 0.4, 0.5\}$, $d = \{0.000075, 0.0001, 0.000125\}$ and $\eta = \{120, 130, 140\}$. For the Full and No Bias models the first search parameters were identical as above, apart from θ being fixed at 0 or 1, respectively. Second search was same for both models with $d = \{0.000075, 0.0001, 0.000125\}$ and $\eta = \{140, 150, 160\}$.

For each set of parameters, 1000 trials were simulated at each possible combination of option values. To evaluate the models using each parameter combination the log-likelihood of each model was computed. The empirical and simulation response times were split into 500ms bins from 1s to 10s. For the simulations the probability of a trial occurring in each time bin was calculated and for the empirical data the amount of trials in each bin was counted. The logarithms of the probabilities were multiplied with the amount of trials in the corresponding empirical time bin and then summed up. The resulting numbers were compared with less negative numbers indicating better fits.

To assess and compare the fit of the final models of each type, likelihood ratio tests were performed using the log-likelihood values calculated as above. The likelihood ratio statistic is calculated as:

$$LR = 2(LL_1 - LL_2)$$

where *LL* denotes the log-likelihood of the models being compared. The likelihood ratio statistic is distributed as $\chi^2_{(1)}$.

Additional analysis

The models were also compared to the empirical data from the even numbered trials on a number of measures. To assess the fit of each model to the empirical data, goodness of fit statistics were calculated. When the dependent variable was binary, proportional tests using χ^2 goodness of fit statistics were employed. When the dependent variable was continuous weighted least squares (WLS), regressions were run on the dependent variable corrected by the empirical average. Weights were equal to the inverse of the empirical variance at that level of the independent variable (usually value-difference). If the models fit the empirical data perfectly the resulting regression models should have zero slope and intercepts. To facilitate qualitative comparison between models, the resulting regression line was integrated with the resulting numbers capturing the deviation from perfect fit (area under curve, AUC).

Results

Overall fits of models

The best fitting parameters for the three models are shown in Table 1. The Free Bias model was compared to the Full Bias and No Bias model using the likelihood ratio test statistic. The Free Bias model's parameters capture the

Table 1: Best fitting parameters for the three models and their corresponding log-likelihood values.

Model	Best fitting parameters			Log-likelihood
	θ	**d**	**σ**	
Full Gaze Bias	0	0.0001	0.015	-3314
No Gaze Bias	1	0.0001	0.015	-3243
Free Gaze Bias (aDDM)	0.4	0.000125	0.01625	-3235

training data significantly better compared to the Full Bias model ($\chi^2_{(1)} = 158, p < 10^{-16}$) and compared to the No Bias model ($\chi^2_{(1)} = 12, p < .001$). This suggests that a model with a moderate to high amount of gaze bias ($\theta = 0.4$) can better account for participants response time distribution compared to similar models which completely discount gaze or allow for full attentional bias during evidence accumulation.

Basic psychometrics

All three models were further evaluated against the even numbered trials in the empirical data to better assess what aspects of the data and the underlying decision process they are able to capture.

First, the relationship between the rated value differences between options and response times was investigated. The empirical data indicated roughly linearly decreasing response times as a function increasing difference between options. When choices are easier, i.e. when the relative value of the options should be easier to discern, choices are also faster (**Fig. 1a**). Comparing the three models shows that all three qualitatively capture this aspect of the data. The Full Bias model performs worst (AUC = 8.91) and consistently overestimates response times. The No Bias model (AUC = 4.08) posits a steeper linear relationship between response times and value difference leading it to overestimate response time for lower differences. The Free Bias model (AUC = 2.72) outperforms the other models in all value difference bins apart from at the highest difference level. These comparative results are, as expected, similar to those from the parameter fitting, which also uses response times.

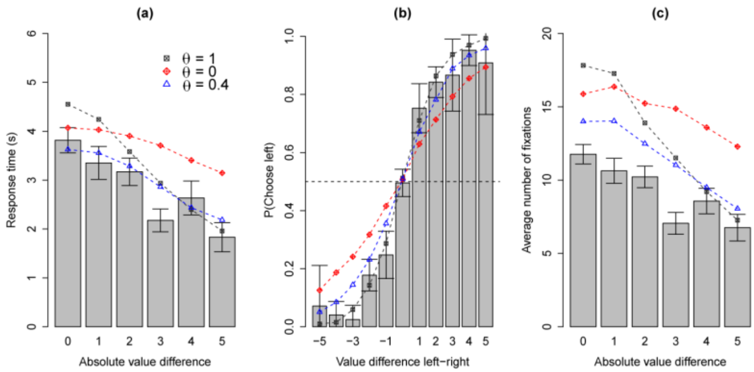

Figure 1: Basic psychometrics. (**a**) Average response times as a function of the value difference between options. (**b**) Probability of choosing the option presented on the left-hand side of the screen as a function of the value difference between the left and right option. (**c**) Average number of fixations in a trial as a function of the value difference between options. Grey bars represent even trials of the empirical data. Error bars denote 95% confidence intervals.

Next the choice curves as a function of the value difference between the options was examined. The empirical data shows a strong sensitivity to value difference with the probability of choosing the lower rated option rapidly decreasing (**Fig. 1b**). Interestingly, the empirical data appears to be asymmetrical with a slight rightward bias in choice when the right option is much higher rated compared to the left option. The models differ in how sharply they predict choice to be guided value; with increasing gaze bias (decreasing θ) the predict choice curves flatten. None of the models predict the empirical data, but the No Bias model provides the best fit ($\chi^2_{(11)}$ = 28.34, p = .003), followed by the Free Bias model ($\chi^2_{(11)}$ = 78.78, $p < 10^{-11}$). The Full Bias model performs the worst of the three ($\chi^2_{(11)}$ = 326.47, $p < 10^{-16}$). This suggests that participants are highly sensitive to the underlying values of the options when computing their final choice.

With increasing value difference fewer fixations are expected. This is also seen in the empirical data (**Fig. 1c**). The Free Bias model performs the best (AUC = 23.77), followed be the No Bias model (AUC = 24.57) and last the Full Bias model (59.49). All models, however, overestimate the number of fixations indicating that the fixation process is not adequately represented in the models.

Exposure and gaze direction

The empirical data shows that, on average, participants are more likely to choose the alternative which they have looked at more (**Fig 2a**). The Free Bias model captures this relationship the best and produces a good fit with the empirical data ($\chi^2_{(11)} = 10.39$, $p = .50$). The Full Bias model also fits the empirical data fairly well, but tends to overestimate the effect of exposure on choice ($\chi^2_{(11)} = 19.41$, $p = .054$). The No Bias model does not predict any relationship between exposure and choice and fits the data poorly in this regard ($\chi^2_{(11)} = 55.31$, $p < 10^{-8}$).

Participants select the last fixated option in 70.2% of trials, but the likelihood of a final fixation being directed at the chosen option does not change with the relative value difference between alternatives (logistic regression, β = -.03, $p = .455$). The Full Bias model fits the empirical data ($\chi^2_{(6)} = 7.41$, $p = .28$), even if it underestimates the amount of fixation bias towards the chosen option. The Free Bias model performs much worse ($\chi^2_{(6)} = 27.80$, $p < .001$), highly

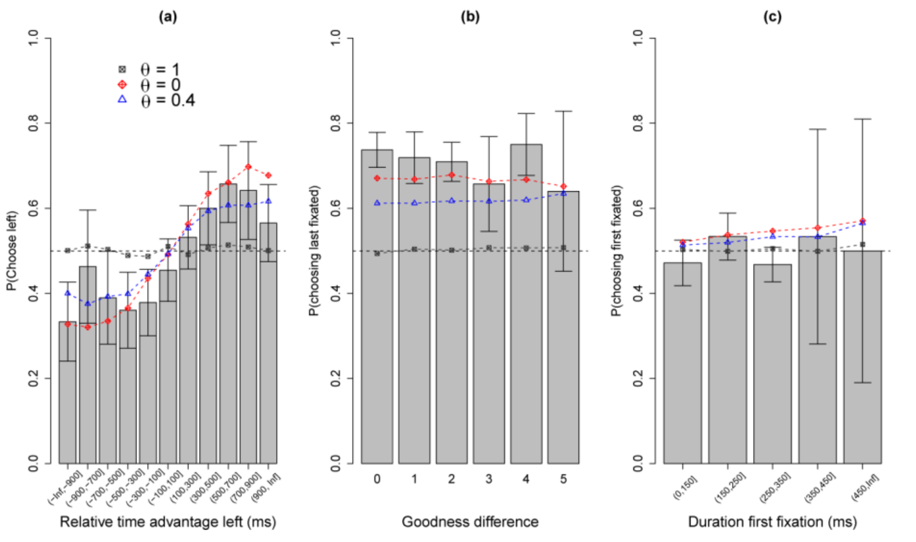

Figure 2: Exposure and gaze direction. (**a**) Probability of choosing the option presented on the left-hand side of the screen as a function of the relative time advantage of that option, in 200ms bins. (**b**) Probability of choosing the last fixated option as a function of absolute value difference between the options. (**c**) Probability of choosing the first fixated option as a function of the duration of the first fixation in a trial. Grey bars represent even trials of the empirical data. Error bars denote 95% confidence intervals.

underestimating the proportion of trials terminating with fixations directed towards the chosen option. Since it posits no relationship between gaze and choice, the No Bias model cannot capture this aspect of the empirical data ($\chi^2_{(6)} = 124.25, p < 10^{-16}$).

Lastly, on a gaze-dependent model, the duration of the first fixation should bias choice. The empirical data does not show this relationship (**Fig. 2c**), and consequently the No Bias model fits the data best on this measure ($\chi^2_{(5)} = 2.09, p = .84$). The Free Bias model predicts a lower amount of first fixation bias ($\chi^2_{(5)} = 4.18, p = .52$), producing a better fit than the Full Bias model ($\chi^2_{(5)} = 5.73, p = .33$).

Discussion

To summarise, three versions of the attentional drift diffusion model (aDDM) were fit to data on charitable choices. The Free Bias model, which was characterised partial gaze bias, proved a better fit to the data compared to the alternative Full and No Bias models. These results mirror existing findings reported in the literature, both for economic choices (Krajbich et al., 2010), as well as for moral choices (Pärnamets, Balkenius & Richardson, 2014). This provides further evidence that attention-mediated decision mechanisms operate at all levels of human choice.

There are several important aspects to the findings presented here. First, the models tested indicate that the aDDM can be used to fit response time data from decisions about charitable choices. This expands the scope of applicability of diffusion models as a general framework for organising data about human decision-making. Second, the aDDM, as tested here, can account for other parts of the empirical data, such as the empirical choice distributions. The model also provides mechanistic accounts for how the direction of the last fixation and exposure differences between options can bias the decisions studied. Third, the results presented here corroborate previous work on eye gaze and moral choice that have showed that using purely text-based stimuli to represent abstract moral alternatives is sufficient for gaze-dependent choice mechanisms to operate (cf. Pärnamets et al., in press). Together, this highlights the fundamental embodiment of human cognition and the pervasive coupling between eye gaze and attentional mechanisms during choice.

To improve future modelling and build a better understanding of the processes underlying choice, it is important to consider not only the aspects of the empirical data where the Free Bias model provided good fits, but also where the two alternative models performed better. The first question concerns how fixations are modelled. As is evident from Fig. 1c, the current method of generating fixations overestimates the number of fixations. Given the relatively accurate response time predictions, this implies that the current method of generating fixation durations results in too many fixations of short durations. It might be that, in the empirical data, participants' fixation durations are particularly associated with some portion of the decision process, such as shorter fixations early on in a trial and longer fixations later on. It might also be a result of reading behaviour early on when participants might be orienting between the decision alternatives.

The question concerning fixations connects to another issue for accurately modelling fixations and gaze in the context of a decision task; how transitions between options are represented. In the current model, transitions are simply represented as a Markov chain. Recent work on consumer decisions has, instead, used saliency as a key component driving fixations, while remaining within an aDDM model, thereby improving the predictive power and, arguably, external validity of the model (Towal, Mormann & Koch, 2013). However, in the present context, such a route seems closed, as the current stimuli consist of black type on grey background. Hence, while there are promising suggestions that information gain might drive the 'why' of transitions (cf. Gottlieb & Balan, 2010), it is not clear how to operationalise this to model the 'when' of transitions for the material used here. One interim solution might be to consider another functional unit instead of fixations as the basis for analysing choices between highly abstract stimuli, like the ones presently studied. For example, it could be possible to use dwells as the fundamental unit, rather than a series of fixations. how dwell durations might affect choice is an unexplored question, especially with the longer decision times found for moral choices compared to simple economics ones.

An interesting contrast to previous work on economic choice can be found when considering the findings represented in Fig. 2. For each of the three analyses presented there (a-c), each of the three models outperformed the others once. This demonstrates that the overall capacity of the Free Bias model to account for the full range of the empirical data is not maximal. Considering, first, the effect of the duration of the first fixation (Fig. 2c) it is possible that, as discuss above, first fixation is the simply the wrong functional unit to analyse. However, the patterns of the empirical data might also reflect discounting of early evidence, and the existence of leaky integrators in the accumulation process (Usher & McClelland, 2001).

The possiblity that leaky integration underpins evidence accumulation in this task is worth considering in greater detail, because it could also explain why the Full Bias model better predicts the empirical data with regard to the role of the last fixation on choice (Fig. 2b; cf. Pärnamets et al., in press). It appears as if participants, in the final moments leading up to their choice, behave as if they are fully discounting the value of the non-fixated alternative. Another alternative, staying within the aDDM framework, would be to consider a hybrid model with θ perhaps shifting over the course of a trial. Finally, with regard to Fig. 2a, an important follow up for future work is to investigate to what extent the effects of exposure found are merely the result of effects of final gaze direction, or if there is a gradual build up of exposure difference during the course of a trial.

So far little has been said about the implications of these findings for understanding morality in general. Providing a computational framework within which moral cognition can be understood is an important development. Not only does it offer a route towards theoretical integration with other decisions, but it can also inform our understanding of moral decisions directly. One aspect is the discovery is that eye gaze is part of the mechanism underlying moral choice. This ultimately means that, at times, what we consider to be right or wrong, will depend causally on how we have visually interacted with a morally salient situation. The conditions when this holds and the limits of this relationship, if any, are important targets for future research. Such research becomes more powerful when formulated within an exact framework, like the one used in the present work, as it allows for specific predictions to be tested and evaluated. A second aspect is that on the current framework, moral decisions depend on valuations of alternatives. Therefore, understanding how such valuations arise and how the decision-maker accesses them becomes central to understand her moral choices.

To sum up, the present study adds further evidence that moral decisions are in part reliant on and determined by domain general choice mechanisms (Pärnamets et al., in press; Pärnamets, Balkenius & Richardson, 2014). That is, moral decisions depend on a process of noisy cumulative integration of available options' values that is biased by visual attention and, hence, real time interaction with the immediate task environment.

Acknowledgments: Thanks to R. Bååth for statistical discussions and C. Balkenius, P. Johansson and A. Lind for comments on earlier versions of this paper. This research was generously supported by Stiftelsen Erik och Gurli Hultengrens fond för filosofi vid Lunds Universitet.

References

Armel, K. C., Beaumel, A., & Rangel, A. (2008). Biasing simple choices by manipulating relative visual attention. *Judgment and Decision Making*, *3*(5), 396-403.

Gottlieb, J., & Balan, P. (2010). Attention as a decision in information space. *Trends in Cognitive Sciences*, *14*(6), 240-248.

Hare, T. A., Camerer, C. F., Knoepfle, D. T., O'Doherty, J. P., & Rangel, A. (2010). Value computations in ventral medial prefrontal cortex during charitable decision making incorporate input from regions involved in social cognition. *The Journal of Neuroscience*, *30*(2), 583-590.

Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? *Perception 36 ECVP Abstract Supplement*.

Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, *108*(33), 13852-13857.

Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, *13*(10), 1292-1298.

Krajbich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in Psychology*, *3*:193. doi: 10.3389/fpsyg.2012.00193

Milosavljevic, M., Navalpakkam, V., Koch, C., & Rangel, A. (2012). Relative visual saliency differences induce sizable bias in consumer choice. *Journal of Consumer Psychology*, *22*(1), 67-74.

Moll, J., Krueger, F., Zahn, R., Pardini, M., de Oliveira-Souza, R., & Grafman, J. (2006). Human fronto–mesolimbic networks guide decisions about charitable donation. *Proceedings of the National Academy of Sciences, 103*(42), 15623-15628.

Pärnamets, P., Balkenius, C. & Richardson, D. C. (2014). Modelling moral choice as a diffusion process dependent on visual fixations. In Bello, P., Guarini, M., McShane, M. & Scassellati, B. (eds.) *Proceedings of the 36th Annual Conference of the Cognitive Science Society*. Cognitive Science Society, Austin, TX.

Pärnamets, P., Hall, L., & Johansson, P. (2014). *I see your dilemma: Visual attention and moral choice.* Manuscript submitted for publication.

Pärnamets, P., Johansson, P., Balkenius, C., Hall, L., Spivey, M.J. & Richardson, D.C. (in press). Biasing moral decisions by exploiting the dynamics of eye gaze. *Proceedings of the National Academy of Sciences.*

Shimojo, S., Simion, C., Shimojo, E., & Scheier, C. (2003). Gaze bias both reflects and influences preference. *Nature Neuroscience*, *6*(12), 1317-1322.

Towal, R. B., Mormann, M., & Koch, C. (2013). Simultaneous modeling of visual saliency and value computation improves predictions of economic choice. *Proceedings of the National Academy of Sciences*, *110*(40), E3858-E3867.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, *108*(3), 550-592.

Appendix A - List of charitable organisations

Cancerfonden
Världsnaturfonden WWF
Röda korset
Hjärt- och lungfonden
Barncancerfonden
UNICEF
Rädda Barnen
SOS Barnbyar
Sjöräddningssällskapet
Stadsmissionen
Radiohjälpen
Plan Sverige
Naturskyddsföreningen
Greenpeace
Amnesty
Diakonia
Läkarmissionen
Läkare utan gränser
Barnfonden
Erikshjälpen
Frälsningsarmén
Min stora dag
Hjärnfonden
Astma och allergiförbundet
BRIS
Friluftsfrämjandet
Djurens rätt
Föreningen Fairtrade
Individuell Människohjälp
Hörselskadades riksförbund
Reumatikerförbundet