**Incidental impressions formed from visual social cues;
updating and group context**

Authors: Rachel Newey¹, Kami Koldewyn¹, Richard Ramsey²

¹ Wales Institute for Cognitive Neuroscience, School of Psychology, Bangor University, Bangor, Wales, LL57 2AS, United Kingdom.

² Department of Psychology, Macquarie University, Sydney, NSW 2109, Australia.

Corresponding authors

E-mail: k.koldewyn@bangor.ac.uk / richard.ramsey@mq.edu.au

Link to data and analysis code:

https://osf.io/bx3gj/?view_only=a696123ca1a5443abeb9b94b7a7777c4

Abstract

A variety of social cues, including gaze behaviour, are used to form impressions of others. For example, if another's eye-gaze reliably helps or hinders us while we complete a task, we incidentally form a positive or negative impression about them. In real life, people are rarely so consistent in their behaviour, and they are often encountered in dynamic group contexts. To date, however, it is not yet known how incidental impressions based on eye-gaze cues are affected by either changes in target individual's behaviour over time, or by the group's behaviour. To better understand how impressions are formed when social behaviours change valence over time, we manipulated helping behaviour both at the level of the individual (Experiments 1-3) and the wider group (Experiments 4 & 5). We found no evidence that impressions were driven by initial behaviour (primacy effects). Instead, people tended to form impressions based on the most recent behaviour, with some influence from the overall, average behaviour. In addition, we found that individuals' behaviours appear to be viewed more or less favourably, depending on the behaviour of the wider group. Overall, we demonstrate that impression formation based on eye-gaze cues is not dominated by a single process, but instead reflects a complex product of cognitive mechanisms that integrate average valence over time, the direction of behaviour changes, the recency of observed behaviour, and the group context in which the behaviour is observed.

Keywords: Impression formation; impression updating; gaze-cues; inconsistent information; group context; trust decisions

Introduction

Impressions of others play an important role in social interactions. For example, whether we think someone is trustworthy or deceitful influences how we might interact with them. But since people's social behaviour is dynamic and changes over time, it is important to detect variations in a person's behaviour and update our impressions accordingly. In addition, many social cues are often experienced peripherally to our primary focus of attention, meaning any social inferences we make about a person are incidental to any explicit social goal. For example, during a team meeting with colleagues, we form impressions based on a whole host of social cues that are not task-relevant, including the person's tone of voice, body language and eye-gaze patterns. Although links between social cues and impression formation have been studied previously, the factors that determine how gaze cues that change over time impact incidental impression formation are not well-established. Therefore, the current series of experiments investigates how sensitivity to changing eye-gaze behaviour influences incidentally formed impressions.

To investigate how impressions are updated over time, researchers often explore how initial impressions are updated when new conflicting or inconsistent information is received (Campellone & Kring, 2013; Chang et al., 2010b; Cone et al., 2017; Delgado et al., 2005; Ecker & Rodricks, 2020; Fareri et al., 2012; Mann & Ferguson, 2015; Maurer et al., 2018; Mende-Siedlecki, 2018; Mende-Siedlecki & Todorov, 2016; Rezlescu et al., 2012; Shen et al., 2020; Wyer, 2010; Zarolia et al., 2017). Some studies show that initial impressions continue to influence behaviour when new information is ambiguous, though final impressions are updated (e.g., Delgado et al., 2005; Fareri et al., 2012). Others find that initial appearance (e.g., Rezlescu et al., 2012) or morality (e.g., Zarolia et al., 2017) based indirect impressions continue to exert influence on final impressions when direct experience is available, while others have shown a rapid reversal when new information prompts a re-evaluation of initial information (e.g., Cone et al., 2017; Mann & Ferguson, 2015). These findings demonstrate that in some circumstances (but not all), impressions can be revised following changes in trait-diagnostic behaviour.

A feature of many of these studies is that they involve written descriptions of behaviour, which provides highly salient, valenced, and categorical social information about individuals who are also the focus of attention. For example, observers may witness someone do something unequivocally bad followed by something good (e.g. Mann & Ferguson, 2015). Such designs provide a strong experimental manipulation, while also being quite unlike how most impressions are formed in daily life, which rarely involve such complete reversals of heavily valenced and trait-diagnostic behaviour or the depiction of people as ‘saints’ or ‘sinners’ (Siegel et al., 2018). Indeed, social information that speaks to a person’s character is rarely so unambiguously diagnostic and in the real-world, social behaviour often needs to be tracked over time and “decoded” before it can be used to infer people’s dispositions. Further, impressions are often formed outside of our awareness (Uleman et al., 2008) or more incidentally, as ‘by-products’ of our other goals (Carlson & Mae, 2003), meaning our attention is often focused elsewhere as we process people’s social behaviour.

Instead of using explicit descriptions of behaviour to study impression formation, an alternative approach has made use of more implicit, non-verbal social signals, such as eye-gaze cues. Gaze is a social cue to which humans are particularly sensitive (for reviews, see Argyle & Cook, 1976; Dalmasso et al., 2020; Emery, 2000; Frischen et al., 2007; Hamilton, 2016), which we can use to infer other’s mental states, attentional focus and intentions (Argyle & Cook, 1976; Colombatto et al., 2020; Doherty, 2006; Gobel et al., 2015; Ricciardelli et al., 2002, 2013). Using a gaze-cueing paradigm (Driver et al., 1999; Friesen & Kingstone, 1998; Langton & Bruce, 1999), modelled on a visual attention cueing task (Posner, 1980), task-irrelevant faces that always provide valid gaze-cues (i.e., look toward the correct location) are judged to be more trustworthy than those that always provide invalid gaze-cues (i.e., look toward the incorrect location) (Bayliss & Tipper, 2006; Manssuer et al., 2015; Strachan & Tipper, 2015; Strachan et al., 2017, 2020; Strachan & Tipper, 2017). The effects of gaze-cue behaviour also generalise to impact trust-based decisions in economic games, such that participants risk investing more money with valid (helpful) faces than

invalid (deceitful) faces (Rogers et al., 2014). These findings demonstrate that even when there is no impression formation goal, the observed behaviour is irrelevant to task completion, and the salience of trait-diagnostic signals is relatively low, incidental impressions about an individual's character are nevertheless formed and influence subsequent social behaviour.

Although these initial studies provide strong evidence that we can learn about others' trustworthiness through (task-irrelevant) eye-gaze cues, many questions remain unanswered. For instance, prior work has only used wholly valid or wholly invalid gaze-cueing manipulations. In other words, a given individual always cues either the correct (i.e., always helps) or incorrect (i.e., always deceives) location. In real life, however, people are rarely so consistent in their behaviour and also may change their behaviour to be more (or less) helpful over time, meaning impressions may be harder to form or require updating over time (Brambilla et al., 2019; Mende-Siedlecki, 2018). As such, it is not yet known how less categorical (e.g., helps some, but not all of the time) and evolving (e.g., helps less over time) gaze-cueing behaviour may impact impression formation. In addition, social cognition does not end at the level of the individual; we also process the behaviour of the wider group when forming impressions (Carlston et al., 2015; Cline, 1956; Lammers et al., 2018; Lee & Harris, 2013; Newman & Uleman, 1990; Simpson & Ostrom, 1976), which may impact upon how we perceive individuals within the group. Although the social milieu is likely to be important for understanding how social cues are interpreted in the context of impression formation, to date, eye-gaze cues have not been studied in relation to the social behaviour of the wider group.

The current work, therefore, seeks to probe the factors that contribute to impression formation both when gaze-cue behaviour changes over time and in light of complex group dynamics. More specifically, we investigate how incidental impressions are affected when the helpfulness of individuals' behaviour changes over time (Experiments 1-3) and how changes in group context over time can affect impressions of individuals (Experiments 4 &

5). Understanding the kind of factors that affect impression formation in such dynamic contexts is important, given that incidental impression formation based on unattended social information is ubiquitous and operates in nearly all social exchanges.

General Methods

Transparency & Openness

For each of the main experiments, primary research questions and hypotheses, planned analyses, target sample sizes and exclusion criteria were pre-registered. A link to each pre-registration is provided in its relevant experimental section. Raw data and analysis scripts for each Experiment are available from the project's Open Science Framework page (https://osf.io/bx3gj/?view_only=a696123ca1a5443abeb9b94b7a7777c4). All manipulations and measures are reported, and all exclusions are specified within individual experiments, as well as summarised together (see Supplementary Materials 1).

Replication and Pilot experiments

Prior to conducting the main experiments, we first wanted to replicate a key result from an experiment on which the current work is based (Rogers et al., 2014) to check that our design would result in similarly-sized effects of gaze-cue validity on social impressions, despite our paradigm being slightly different. Rogers and colleagues (2014) used an object categorisation task, where participants identify the category to which the target belongs (for example, kitchen vs garage tools). In these designs, although the eyes provide a cue regarding where a target will appear, once the target is seen the participant must still additionally identify the correct category. Therefore, the judgement is orthogonal to the location of the target. We opted to use a target location version of the task, in which participants simply respond to the location of the target. In this version, although the faces are still incidental to the task (i.e., not necessary to complete the task), they do help (or hinder) participants' performance by identifying (or lying about) the target's future location.

We selected this version to make the faces more genuinely 'social' through displaying directly helpful or deceitful behaviour. This means that while faces' gaze-cues are no longer task irrelevant, impressions formed are still incidental. Happily, despite the change in task, and a few other minor differences (see Supplemental Materials 2 for full details), the replication experiment produced similar between-condition effect sizes, with wholly valid faces being trusted much more than wholly invalid faces (original units £1.33; $t(14)=2.82$, $p=.014$, $CI[0.32, 2.35]$, $dz=.0.73$).

The planned design of our main experiments also differed from prior work by using a range of gaze-cue validities. In other studies, while distractor faces sometimes displayed non-predictive cues (i.e., cues were valid 50% of the time), target faces always provided either only valid cues (i.e., they always looked towards the side where the target would appear) or only invalid cues (i.e., they always shifted their gaze to the side opposite of where the target would appear). Our design required validities that varied both among faces and over time. In addition, gaze-cue validities that vary between 100% and 0% better reflect social behaviour in the real world, where people are rarely fully consistent. We piloted our paradigm to confirm that people are sensitive to a range of gaze-cue-validities and, in general, found that people's impressions can reflect a graded pattern of gaze-cue behaviour (see Supplementary materials 2 for full details). We found a clear and similarly sized difference between wholly valid and wholly invalid faces (original units £2.00; $t(14)=2.69$, $p=.018$, $CI[0.41, 3.59]$, $dz=.69$), together with a clear effect across all six conditions ($f(5,70)=5.37$, $p<.001$, $\eta^2=.28$). This preparatory work provided confidence that the basic experimental design was working as anticipated. Importantly, it also provides evidence that people are sensitive to non-categorical, inconsistent social behaviour and reflect that inconsistency in their final impressions.

Ethical Approval

All experiments were granted ethical approval by the School of Psychology at X University's Ethics Committee, and all participants were X University students who provided

written consent prior to data collection taking place. In return for their time, participants received either course credit or a small cash payment (£4).

Materials & Stimuli

Face stimuli

Six young, white adult male faces, each displaying a neutral expression, were used in all experiments to depict the various behavioural patterns. The faces, together with adapted versions to portray averted gaze to the left or right, were kindly provided by Strachan and Tipper (2016) who used them in a very similar gaze-cueing design and collected baseline ratings of attractiveness and trustworthiness for each face. These six (among others) were matched on both judgements. Faces (IDs: AM 12, 16, 24, 25, 26, & 35) were originally taken from the KDEF face set (Lundqvist et al., 1998).

Gaze-cueing task

A gaze-cueing task was used to portray targets' behaviours (**Figure 1**). Participants were asked to keep one finger from their left hand over the "c" key and one finger from their right hand over the "m" key and their task was to press the correct key when the target appeared; "c" for on the left, "m" for on the right. Participants were instructed to respond as quickly and as accurately as possible when the target appeared and instructions were provided on screen before the task began. There was no mention of the faces. Each trial began with a central fixation cross for 500ms followed by a forward-facing face for 1,500ms, which served to promote joint attention. The face then averted its gaze either to the left or right and after 450ms a letter "T", which served as the target, appeared to the left or right of the face. Once the target appeared, the screen remained the same until the participant made a response or 2,000ms lapsed, whichever came sooner. At this point the target disappeared and the face returned to looking forward for 1,000ms (See **Figure 1**).

Valid trials were those where the gaze-cue direction matched the eventual location of the target. Invalid trials were those where the gaze-cue direction did not match the eventual

location of the target. As such, we refer to valid cues as being helpful because they provide accurate information and invalid cues as being deceitful because they provide inaccurate information. Any change to a face's validity (helpfulness) over time always occurred halfway through the task; between the second and third blocks.

The task was structured into four blocks, with each block including 48 trials, for a total of 192 trials. This resulted in each of the six faces being presented eight times per block and 32 times overall. Each face looked to the left and right equally, and the total number of valid and invalid trials (across all faces) were equal across blocks and overall (with the exception of Experiments 4 and 5). To reduce anticipation and/or carry-over effects, no more than three trials of the same type (i.e., valid/invalid) could occur in succession. Blocks were pre-constructed so that trial order was the same for every participant within a particular block. Blocks 1 and 2 always came before blocks 3 and 4 but the order between blocks 1 and 2 and between blocks 3 and 4 was counter-balanced across participants. Most importantly, which face represented which condition was also counterbalanced across participants, ensuring that any effects of the characteristics (e.g., attractiveness, small differences in facial expression) of particular faces were minimised, as the same face that was always deceitful in the experience of one participant might provide the most help to another. The task took around 15 minutes to complete. To allow participants to rest their

Figure 1: An example of an invalid gaze-cueing trial, where the target appears on the side opposite to the location gazed at by the cue provider (ID: AM24)

eyes, there was a break between each block where the task paused until the participant chose to continue. Analyses for gaze-cue task data for each experiment can be found in Supplementary Materials 3.

Trust

Our primary dependent measure was taken from the trust (investment) game (Berg et al., 1995). In the 'one-shot' variant used here, participants had hypothetical £10s and were allowed to invest any amount from £0 to £10 with each trustee (face) only once. Participants were informed that any amount invested would be multiplied by three and the trustee would then have full discretion as to how much, if any, they returned to the participant. In order to check that participants understood the task and had the opportunity to ask questions, they completed two practice trials where they were shown two male faces (not included in the gaze-cueing task) and asked to decide how much they would invest with each one prior to the computerised task. Given the hypothetical nature of the task, participants were asked to act as though the decisions they were making had real consequences and would actually result in them winning or losing money. In this version of

the game, often referred to as a 'one-shot' trust game, only the first 'investment' step was used and the trustees' responses were not revealed.

Face ratings

On a computer screen, participants were shown all the faces that they had seen during the gaze-cueing task and were asked to rate how nice they thought each face was using a 7-point Likert scale, where 1 represented "not nice at all" and 7 represented "very nice". As our primary measure of interest concerned trusting decisions, analyses for the niceness ratings are reported in supplementary materials. We observed smaller, yet similar, effects across trusting decisions and niceness ratings. Analyses for niceness data for each experiment can be found in Supplementary Materials 4.

Procedure

Upon arrival, participants were informed they were going to complete three tasks; a reaction time-based computer task, a face rating computer task and some questionnaires. Included questionnaires varied across experiments and, as they did not form part of any primary hypotheses, their data is not presented here. First, the trust game was explained, where participants were given an instruction sheet and the premise was re-iterated by the experimenter before the two practice trials were completed. Second, the computerised tasks were completed, starting with the gaze-cueing task, which was described to participants as a reaction time task. Immediately after, each face presented during the gaze-cueing task was rated on niceness ("How nice do you think this person is?") followed by the one-shot trust games ("how much would you like to invest with this person?"). To reduce anchoring effects, the order in which faces associated with each condition were presented was randomised across participants.

Before being debriefed, participants were asked if they had guessed the purpose of the experiment. While many said that the experiment was "looking at how we judge people's faces", none made the direct link between the behaviour of the faces during the gaze-cueing

task and its effect on judgements. Indeed, none reported noticing that the faces differed in their behaviour or that some faces were 'helping' while others were not. The whole experiment, including ratings, took approximately 30 minutes to complete.

Data Analysis

Prior to analysis, responders designated as 'response outliers' were removed from the datasets. A participant was pre-registered as a response outlier if they responded with the same likeness rating or investment amount for all six faces presented. This means that a participant could be removed from niceness analyses but still be included in trust analyses. While we cannot be sure that repeated responses were not genuine, these data were not included in final analyses as they are more likely due to a participant not engaging with the task. The number of response outliers was very low (range: 0-3) across all experiments (see Supplementary Materials 1). We also analysed the data without the removal of outliers and it made no meaningful impact on the primary results.

Though the results of the gaze-cueing tasks themselves were not part of the main question, we analysed the gaze-cueing data to ensure that the task generated the expected gaze-cueing effect; namely that invalid gaze-cues generated longer reaction times than valid trials, and that overall error rates were low. Incorrect trials and those with reaction times of less than 200ms or greater than 1,500ms were removed prior to all analyses. Average accuracy and reaction time data was then calculated based on trial validity and block number. Participants who scored less than 80% overall accuracy were designated as 'gaze-cueing' outliers and were removed from all analyses. Though this was not pre-registered, accuracy rates are generally high for gaze-cueing tasks, so a low accuracy score is suggestive of poor attention and, for our purposes, means it is less likely that the faces' behaviours were implicitly encoded. The number of gaze-cueing outliers was very low (range: 0-5) across all experiments (see Supplementary Materials 1). Analyses of all Experiments' gaze-cue data can be found in Supplementary Materials 3.

Following the removal of outliers, our general approach was to first use a RM ANOVA and then follow with paired contrasts using mean differences (effect sizes in original units) and 95% CIs. We also report standardised effect sizes using Cohen's d_z calculated as $t/\sqrt{n}$. More specific details on analyses are provided in each experimental section. Where violations of sphericity occurred, Greenhouse-Geisser adjusted values are reported.

Following proposals by Gigerenzer (2018), we avoid interpreting results based solely on p-values and binary distinctions between “significant” and “non-significant”. Instead, we base the direction and strength of our inferences on a range of metrics, which include a mixture of descriptive and inferential statistics. Hypothesis tests and p-values are considered in combination with an estimation approach, which emphasises the importance of estimating effect sizes (in original and standardised units) along with a measure of precision using 95% CIs (Cumming, 2014). Further, we have embedded conceptual replication into the design of the Experiment by using the same basic experimental design over a series of experiments (Zwaan et al., 2018).

Experiment 1

Introduction

This experiment investigated how inconsistent gaze-cue behaviour over time and the order in which helping behaviour is provided influence impression formation. Gaze-cue behaviour was manipulated such that individual faces either increased, decreased or maintained their proportion of valid (helpful) cues over time, while also creating two distinct groups of faces where overall (average) helping behaviour was the same (i.e., three ‘helpful’ faces vs three ‘unhelpful’ faces) (see ‘Design’ below).

Based on both previous work and our pilot data (Supplementary Methods 2), we expected participants would form impressions of the faces during the first half of the task (blocks 1-2). Here, our primary research question probes how changes in helping behaviour during the second half of the task (blocks 3-4) would affect overall impressions. Impressions

may not be affected and may continue to reflect early behaviour (primacy); they may be updated to reflect later behaviour (recency); or they may represent an assimilation of information across the whole task (averaging). In the current experiment, where participants are experiencing social behaviour that (incidentally) helps or hinders them in completing a task quickly, we expect the most recent behaviour to have a greater impact on final impressions than initial behaviour (Campellone & Kring, 2013; Chang et al., 2010a; Maurer et al., 2018). However, one consequence of providing all behavioural information in one uninterrupted interaction may mean that participants calculate a running average of social information as it is experienced and adopt an ‘anchor and adjust’ (Hogarth & Einhorn, 1992) or average-based approach (Anderson, 1965). A simple average would seem unlikely, however, as negative behaviours (i.e., deceitful gaze-cues) are usually more salient and heavily weighted than positive (Baumeister et al., 2001; Fiske, 1980; Skowronski & Carlston, 1989). For example, when people learn about others’ past (im)moral behaviours, updating is greater for moral-to-immoral behavioural changes than immoral-to-moral changes (Mende-Siedlecki, 2018; Mende-Siedlecki, Cai, et al., 2013; Reeder & Coover, 1986). The direction of change (i.e., helping-to-hindering vs hindering-to-helping) could, therefore, impact the size of differences in perception between changing and stable pairs of faces. Our design helps us probe the influence of initial vs. recent information, the direction of change (helping to hindering or vice-versa), and the extent to which participants are simply averaging their experience over time.

Method

Power analysis

In acknowledgment of the fact that effect sizes are likely to be smaller when reproduced (Camerer et al., 2018; Open, 2015) and given that we are not only seeking to assess a range of validities, but also to see if there is an effect of the order in which they are witnessed (i.e. changing validities over time), we powered the experiment to detect effects smaller than those returned from pilot data (Supplementary Materials 2). Using G*Power 3

(Faul et al., 2007), we targeted the main effect of order (i.e., comparing faces who decreased vs. maintained vs. increased helping behaviour over time) and ran an a priori sample size calculation for a repeated measures design with three conditions. A Cohen's f of .33 ($np^2=.10$) was set with 80% power and an alpha of .05. This returned a requisite sample size of 46, which we rounded to 50. To take account of potential outliers, we aimed to collect 55 datasets (link to pre-registration: <https://aspredicted.org/blind.php?x=fs3mb2>). A sample size of 50 provides 80% power to detect (unadjusted) differences between pairs of conditions of $dz=0.40$ and higher.

Participants

We recruited 55 (female 44, missing 1; mean age 21.8 years, SD 3.9) participants, each of whom was compensated with course credits.

Design

The proportion of total valid and invalid gaze-cues provided by each face was manipulated such that there were two groups comprised of three faces (Table 1). The individuals in one group each provided 75% valid cues (mostly helpful) when averaged across the entire task, while the individuals in the other group each averaged only 25% (mostly deceitful). While all faces within a group had the same overall number of valid cues, the order in which they displayed their helping behaviour varied. One decreased its valid number of cues over time, one increased valid cues and one maintained a consistent level of helping behaviour over time. This resulted in six unique gaze-cue profiles.

Data analysis

Levels of trust were analysed using a 2 (helping group: 75% average help vs 25% average help) by 3 (order: decrease, stable, increase) repeated measures Anova.

Validity %	Overall Helpful Group (75%)			Overall Unhelpful Group (25%)		
	Decrease	Stable	Increase	Decrease	Stable	Increase
Initial (Blocks 1&2)	100	75	50	50	25	0
Final (Blocks 3&4)	50	75	100	0	25	50
Average Validity	75	75	75	25	25	25

Table 1: Gaze-cueing task conditions for Experiment 1.

Of the 55 participants tested, three response outliers were removed from the trust analysis and one from the niceness data, with only one participant being removed from both. Three additional participants were removed as gaze-cueing outliers, resulting in 52 participants being included in the gaze-cueing analyses (Supplementary Materials 3), 49 in the trust analyses and 51 in the niceness analyses (Supplementary Materials 4).

Results

There was a main effect of group $f(1,48)=19.11$, $p<.001$, $np^2=.29$ on investment decisions, with the mostly deceitful group (averaging 25% help) receiving lower (M £3.44, SD 1.44, CI[3.0,3.9]) investments than the mostly helpful group (averaging 75% help) (M £4.50, SD 1.36, CI[4.1,4.9]). There was also a main effect of order $f(2,96)=7.11$, $p=.001$, $np^2=.13$. Main effects were partly qualified by a marginal group*order interaction effect $f(2,85)=2.83$, $p=.071$, $np^2=.06$. To unpack the direction of these latter effects, we visualise the distribution of responses per condition and compute differences between conditions within each group (helpful and unhelpful).

As can be seen in **Figure 2**, across both groups, investments were lower in the decreasing conditions than the other conditions. Sidak adjusted comparisons showed the decrease conditions received less than both the stable conditions (mean difference £0.99, CI[-1.7,-0.3], $p=.005$) and the increase conditions (mean difference £0.79, CI[-1.4,-0.2], $p=.011$). On average, faces that maintained their helping behaviour throughout the task were trusted no more or less than faces that increased their levels of help (mean difference £0.20, CI[-0.5,0.9], $p=.862$), though this was not the same for both groups. As can be seen in

Figure 2: Violin plot for Trust decisions (investments) for Experiment 1. On the left is the helpful group and, on the right, the deceitful group. Bars represent 95% CI.

Figure 2, there was a clear difference between how the conditions in the helpful group were trusted $f(2,96)=8.00$, $p=.001$, $np^2=.14$, but the difference was less apparent for the unhelpful group $f(2,96)=2.45$, $p=.092$, $np^2=.05$. Adjusted comparisons (see **Table 2**) revealed no difference between investments made with the decrease and stable faces in the mostly unhelpful group but a clear difference in the helpful group. There was also a slight difference between the stable and increase faces in the helpful group. In other words, stability was valued most in the helpful group but not in the unhelpful group.

Our original design aimed to explore averaging and order effects. However, some of the results could, at least partly, be attributable to recency effects. That is because in this design, the most helpful faces also ended with higher levels of helping behaviour than the deceitful faces. Likewise, the decrease faces also ended with lower levels of helping behaviour than the other faces. In order to assess whether recency might also partly explain our effects, we directly compared the two conditions that were in different groups and changed their behaviour in opposite directions (one decreasing, the other increasing) but who both ended by helping the same amount (50% validity; conditions 1 and 6, Figure 2). Despite their behavioural differences, there was no difference in trust based decisions.

	Mean difference	95% CI Lower	95% CI Upper	t	dz	p
helpful decrease > helpful stable	-1.49	-2.51	-0.47	3.63	.52	.002
helpful decrease > helpful increase	-0.76	-1.64	0.13	2.12	.30	.114
helpful stable > helpful increase	0.74	0.12	1.59	2.12	.30	.114
unhelpful decrease > unhelpful stable	-0.49	-1.55	0.57	1.15	.16	.589
unhelpful decrease > unhelpful increase	-0.84	-1.66	-0.13	2.51	.36	.045
unhelpful stable > unhelpful increase	-0.35	-1.27	0.58	0.93	.13	.736
helpful decrease > unhelpful increase *	-0.08	-1.15	1.00	0.23	.03	>.99
helpful increase > unhelpful decrease ^	1.51	0.13	2.89	3.38	.48	.022

Table 2: Sidak adjusted pairwise differences in trust decisions for Experiment 1. Note: * these conditions both ended with 50% validity (i.e., recency driven) and ^ these conditions both started with 50% validity (i.e., not primacy driven).

Conversely, there was no evidence for a primacy effect, as the two conditions that started helping the same (50%) amount were not trusted the same (**Table 2**). This suggests that recency may have played a role in determining trust levels in addition to average valence and order effects. Recency alone, however, cannot explain all the results. In the helpful group it is the consistent face, and not the face that ends by always helping, who receives the highest investment. Likewise, in both groups, participants made similar investments to faces that did not end by offering the same level of help.

Discussion

Experiment 1 demonstrates that overall trust levels appear to be the product of a cognitive system that can track the order and recency of behaviour, as well as integrate behaviours into an overall average. Rather strikingly, in our view, even in a social situation where the trait-diagnostic information is subtle and impressions are formed while performing another task, we show that character impressions are formed in a manner that integrates a

range of diverse signals over time. We show no evidence that the initial behaviour displayed by an individual influences trust. In the context of eye-gaze cues that change over time, therefore, we show no evidence of a primacy effect on impression formation. However, the design of Experiment 1 might have been too complicated for people to sufficiently distinguish between the six individual faces over time. Failure to adequately distinguish between behaviours may have reduced the ability of participants to track changes. Therefore, in two subsequent experiments, we used simplified experimental designs to further probe the relationship between social cues and trust.

Experiment 2

Introduction

Experiment 2 was designed to probe the effects of recency in a more direct manner than was possible in Experiment 1, where recency effects were confounded with order effects. We simplified the design two ways. Experiment 2, therefore, used a simplified design with only three conditions, each with a different average validity and sequence of helping behaviour. Crucially, all faces ended by offering the same level of helping behaviour (Table 3). If, despite all ending the same, faces are viewed differently, it would suggest that faces' behaviour is processed and incorporated over time and not that only the most recent behaviour determines overall impressions. In this instance, the face that decreases how much it helps over time may be perceived the most (as it averages the most helpful behaviour) or the least (as it decreases how often it helps over time) favourably.

Method

Power Analysis

As with Experiment 1, a power analysis was performed for a 1*3 repeated measures analysis of variance. Again, the design required a minimum of 46 participants and in consideration of outliers, 55 were targeted (link to pre-registration: <https://aspredicted.org/blind.php?x=q5wq2w>).

Participants

Fifty-two new volunteers took part in this Experiment (41 female; mean age 21.59, SD 4.84), each of whom was compensated with either course credit or a small cash payment (£4).

Design

The validity of the gaze-cues provided by the faces was manipulated such that there were three conditions and only ever three variations in cue validity in any one block. Each condition was portrayed by two faces. As with Experiment 1, faces in this experiment either decreased, increased or maintained their level of helping behaviour over time. In order to disambiguate recency effects from the order and averaging effects seen in Experiment 1, the difference here was that all faces ended by displaying the same pattern of gaze-cue behaviour (*Table 3*).

Data Analysis

Prior to the main analyses, data from the pairs of faces producing the same behaviours were compared using within subject t-tests. There were no significant differences between any of the pairs (all p 's $>.05$, unadjusted), and they were collapsed to create the three conditions; decrease, stable and increase. The main test of interest was a 1*3 repeated measures Anova.

Of the 52 people tested, one person was removed as a response outlier from both the niceness and trust datasets. There were no gaze-cueing task outliers. This resulted in 52 participants being included in the gaze-cueing analyses (Supplementary Materials 3) and 51

Validity %	Decrease	Stable	Increase
Initial (Blocks 1&2)	100	50	0
Final (Blocks 3&4)	50	50	50
Validity %	75	50	25

Table 3: Gaze-cueing conditions for Experiment 2.

in both the niceness (Supplementary Materials 4) and trust analyses.

Results

Visual inspection of the data in

Figure 3 shows that the condition that decreased helping over time, but helped the most overall, had slightly higher investments than the condition that increased how much it helped, but overall provided the least amount of help. A one-way repeated measures Anova, though, revealed no differences in trusting decisions between conditions $f(2,100) = 1.59$, $p = .209$, $\eta^2 = .03$ and an adjusted comparison showed that the difference between the two conditions was too small to provide clear or convincing evidence for a difference (£0.60, CI[-0.3, 1.5], $p = .271$, $d_z = 0.23$).

Discussion

When the task was simplified through the use of only three conditions and all faces ended with the same amount of helping behaviour (50%), there were no clear differences between conditions. The results, therefore, are not consistent with an averaging approach over time being dominant. Indeed, there was only a relatively small difference in trust between the two conditions that changed their helping behaviour over time and offered substantially different levels of helping behaviour overall (75% vs 25%). In this design, where all faces end by offering the same amount of helping behaviour, the results provide support for a recency effect with the potential for only a small influence from overall valence. However, it is possible that the change in behaviour between blocks 2 and 3 may not have been salient enough to strongly impact levels of trust. Therefore, the dominance of recency that we observed here may, in part, reflect the way the experiment was constructed. Experiment 3 was designed to address this issue.

Figure 3: Violin plots for trust decisions (investments) for Experiment 2. Bars represent 95% CI.

Experiment 3

Introduction

Experiment 3 used a simpler design than Experiment 1, while also making the change in behavioural valence between blocks 2 and 3 more distinct than Experiment 2. We aimed to identify if there were circumstances in which primacy, averaging or order effects might better explain impressions than recency effects. If, despite very different initial and average helping profiles, faces that end the same are trusted to the same degree, we would provide much stronger evidence that the most recent pattern of helping behaviour is the most dominant predictor of how impressions are formed. If, however, we find that there are other clear effects present, it would suggest that primacy, averaging or order effects can also contribute to the construction of impressions formed from inconsistent behaviour over time.

Method

Power Analysis

The a priori power calculation for this experiment was based on a 1*5 repeated measures Anova. The previous two experiments were powered to detect medium to large effects and above. As effects detected in Experiment 2 were below that level, this

experiment was powered to detect medium effects and above. At 80% power, with an alpha of .05 and Cohen's f of 0.25 ($np^2=.06$), this returned a requisite sample size of 50. In consideration of potential outliers, 55 datasets were targeted. Differences between specific combinations of conditions were also pre-registered (link to pre-registration: <https://aspredicted.org/blind.php?x=mn54eq>). For paired sample t-tests (2 tailed), 50 participants provide 80% power to detect Cohen's d_z of 0.40 and above.

Participants

Fifty-eight volunteers took part in Experiment 3 (44 female; mean age 22.4, SD 3.1), each of whom was compensated with either course credit or a small cash payment (£4).

Design

The previous two experiments involved changes of 50% in helping behaviour. Here, the changes that occurred were a full 100% reversal in helping behaviour. We also included three stable conditions, whereby individuals remained constant throughout at 100%, 50% or 0% validity (**Table 4**). As a consequence, there were two pairs of conditions that ended by offering the same pattern of helping behaviour, but started by offering opposite levels of helping behaviour, allowing us to compare recency to averaging effects. There were also two individuals that offered the same, non-predictive behaviour throughout, which were collapsed into one condition and formed the control. By only having two faces change their behaviour over time (rather than four as in both Experiment 1 and 2), we hoped that the difference between the faces' behaviour would be more striking.

Validity	Stable	Decrease	Stable (Control)	Increase	Stable
Initial (Blocks 1&2) %	100	100	50	0	0
Final (Blocks 3&4) %	100	0	50	100	0
Average %	100	50	50	50	0

Table 4: Gaze cueing conditions for Experiment 3.

Data Analysis

Data from the two control faces were compared prior to the main analyses. There was no significant difference between the two ($p > .90$) and so they were collapsed, creating five overall conditions. The test used to compare conditions overall was a 1*5 repeated measures Anova. The planned, main tests of interest were paired samples t-tests targeted to compare specific pairs of conditions; the two conditions ending at 100% validity and the two conditions ending at 0% validity. As paired comparisons were planned, we used the least squares difference adjustment for multiple comparisons.

All 58 participants provided some variation in their responses, meaning none were removed as response outliers. Following a review of the gaze cueing data, however, four participants were removed from further analysis; three as their overall accuracy was less than 80% and one because of a technical error (see Supplementary Materials 1 for full details of how data was processed). This resulted in 54 participants being included in the final gaze-cue task (Supplementary Materials 3), niceness ratings (Supplementary Materials 4) and investment decisions analyses.

Results

A repeated measures Anova showed a main effect of condition $f(3, 163) = 6.59$, $p < .001$, $\eta^2 = .11$ (**Figure 1**). To directly address our research questions, we estimated the size of the differences between conditions by performing pairwise comparisons (Table 5). To assess the relative influence of primacy vs. recency, we used paired-sample t-tests to contrast the two pairs that started the same but ended differently (**Table 5**, 'primacy' contrasts) and the two pairs that started differently but ended as the same (**Table 5**, 'recency' contrasts). In line with the results from experiments 1 and 2, there was no evidence for primacy. If the first information was driving impression formation, we would expect to find no differences between those that started the same, but ended differently. Instead, both the

Figure 4: Violin plot for Investment amounts for Experiment 3 where conditions show initial and final gaze-cue validity. Bars represent 95% CI.

contrast between the 100-100 and the 100-0 conditions and the contrast between the 0-100 and 0-0 conditions revealed reliable differences (**Table 5**, ‘primacy’ contrasts).

On the other hand, if the most recent information was driving impression formation, we would expect no difference between those conditions that ended identically, regardless of how they started. Both the contrast between the 100-100 and 0-100, and the 100-0 and 0-0 conditions revealed only small differences, with effect sizes smaller than we can reliably detect (see Table 5, ‘recency’ contrasts). However, these differences do suggest that recency is not the only influence on impression formation. This is particularly evident when considered in “real” terms. Participants on average invested £0.76 more with the wholly helpful face than the face that only helped at the end, a difference that is nearly 8% of the possible investment and nearly 17% of the average investment.

Similarly, we assessed the influence of averaging most simply by comparing the three condition combinations that averaged at 50% validity (100-0, 0-100, 50-50) (**Table 5**, ‘averaging’ contrasts). For the two changing conditions, the face that ended by helping was numerically trusted more than the face that changed to offer no help, but this effect was too small for us to reliably detect. The control face looked to be trusted less than the increasing

Contrast	Mean difference	95% CI Lower	95% CI Upper	t	dz	p
Primacy						
100-100 > 100-0 ^	1.35	0.48	2.23	3.10	.41	.003 *
0-100 > 0-0 ^	0.96	0.19	1.73	2.51	.34	.015 *
Recency						
100-100 > 0-100 ^	0.76	-0.13	1.64	1.72	.23	.091
100-0 > 0-0 ^	0.37	-0.25	1.00	1.18	.16	.242
Averaging						
100-0 > 50-50	0.20	-0.66	1.07	0.69	.09	.999
0-100 > 50-50	0.80	-0.23	1.83	2.26	.31	.248
100-0 > 0-100	-0.59	-1.71	0.53	1.55	.21	.747
Other						
100-100 > 50-50	1.56	0.31	2.81	3.63	.49	.006 *
100-100 > 0-0	1.72	0.25	3.19	3.42	.47	.012 *
0-0 > 50-50	-0.17	--0.98	0.64	0.60	.08	.999

Table 5: Adjusted pairwise comparisons for Experiment 3. ^ these contrasts were planned and were not adjusted

(0-100) face (-80p), but very similarly to the decreasing (100-0) face (-20p). Indeed, the control and decrease conditions are trusted no more than the wholly deceitful (0-0) condition (Table 5, 'other' contrasts). This finding does not support averaging, recency, or order processes, but suggests that individuals who are neutral, or non-predictive, are viewed as negatively as those who are wholly invalid. Nonetheless, it is suggestive that a simple averaging mechanism alone cannot explain the current results.

Discussion

In Experiment 3, when using a design that was simpler than Experiment 1 and provided more distinct changes in behaviour than Experiment 2, the results cannot be explained according to simple recency or averaging explanations alone. Rather, while recent

behaviour is the best predictor of final impressions, the results suggest that encoding social information over time is nuanced and impression formation processes also depend on the overall valence of behaviour *and* the order in which it is provided. Indeed, unhelpful recent behaviour appears to over-write all previously helpful behaviour, but previous unhelpful behaviour may not be fully over-written by subsequent helpfulness. In other words, to be perceived positively, one must be consistently helpful, while negative behaviour impacts impression whether it is recent or not.

Across Experiments 1-3, we focused on an individual's social behaviour (gaze cueing behaviour). However, it is likely that impressions of one face have been influenced by impressions of other faces appearing in the same group (Carlston et al., 2015; Lammers et al., 2018). It would seem reasonable, therefore, that impressions formed from gaze-cue behaviour are relative and a product of both the behaviour of individuals and that of their associated group members. Experiments 4 and 5 test this assertion by presenting the same two behavioural profiles in two different social environments (link to pre-registration: <https://aspredicted.org/blind.php?x=j8y5i4>).

Experiment 4

Introduction

Experiment 4 investigated how consistent individuals are perceived when encountered in two directly opposed changing social environments. The valence of two target individuals' behaviour remained constant over time and across environments, but the behaviour of other group members changed both over time and across environments. In one environment, the overall levels of helping behaviour across group members improved over time, while in the other, group members' behaviour became less helpful over time. We predicted a contrast effect (e.g., Cline, 1956), whereby the targeted individuals would be trusted more in the environment where others become less helpful, rather than when they become more helpful.

Method

Power analysis

In keeping with previous experiments, a sample size of 55 was targeted. These designs were, however, 2*2 repeated measures designs, meaning we are only powered to detect larger effect sizes than those previously targeted. A sensitivity analysis for a two-level repeated measures Anova with an alpha of .05 calculated that partial eta squares of 0.14 ($d_z=0.38$) or higher can be interpreted with 80% power.

Participants

A total of 64 (49 female, 10 male, five unknown; mean age 21.14, SD 5.21) volunteers began this two part experiment, although not all completed both sessions (see data analysis section). Demographic information was collected in the second session of the experiment, meaning that the information is not available for those participants who did not attend the second session. All participants were compensated for their time with course credit.

Design

Two faces maintained a consistent proportion of valid cues across two different environments. One individual was consistently helpful (75% valid) and the other was consistently unhelpful (25% valid). They were encountered in two different social environments, where faces either increased (improving environment) or decreased (worsening environment) how much they helped over time. The changeable faces that made up each environment were mirror images of each other, providing the same amount of help overall, but reversing the order in which they provided that help according to the environment (see **Table 6**). This created a situation where the ratio of valid to invalid cues differed both within and across blocks.

To be consistent with previous experiments, the same six faces were included in the gaze-cueing task, with the changeable conditions being displayed by four faces. The mathematics of the design necessitated the use of two faces to each portray the same condition in each environment (see **Table 6**). To maintain the average levels of helping behaviour displayed by the changing faces in each environment, the middle condition of each was selected for repetition. This was a within-subjects design with the same participants experiencing both environments.

Procedure

The experiment was conducted across two sessions, with a period of at least three days between session one and session two. In the first session, the trust game was explained and practiced. Participants then completed either the improving or the worsening condition before rating each face for niceness and engaging in one-shot trust games. In the second session, participants completed the remaining condition and some questionnaires. The order in which the two tasks were completed was counterbalanced across participants.

Data Analysis

The two stable faces were compared across the two environments using a 2 (Environment: Improving, Worsening) * 2 (Person: Helpful, Unhelpful) repeated measures Anova. To explore face and environmental effects separately, a series of planned paired

Validity	Unhelpful Target	Helpful Target	Improve	Improve*	Improve
Initial (blocks 1&2)	25	75	0	25	50
End (block 3&4)	<u>25</u>	<u>75</u>	<u>50</u>	<u>75</u>	<u>100</u>
Average	25	75	25	50	75
			Worsen	Worsen*	Worsen
Initial (blocks 1&2)	25	75	50	75	100
End (block 3&4)	<u>25</u>	<u>75</u>	<u>0</u>	<u>25</u>	<u>50</u>
Average	25	75	25	50	75

Table 6: The tables show the conditions and validities for Experiment 4 for the improving (top) and the worsening (bottom) environments. *This same condition was displayed by two faces.

samples t-tests were performed.

Due to the two-session nature of the experiment, there was some attrition, with 61 participants completing the improving environment, 59 the worsening environment and a total of 56 completing both. There were no response outliers. There was one gaze-cueing outlier and there were technical issues that resulted in data loss for four participants. These five participants' data was removed from all analyses, resulting in 51 subjects being carried through for final social judgement analyses (see Supplementary Materials 4 for Niceness analyses). Gaze-cue data were analysed separately for each task. In the improving environment, there was one outlier and missing data for two people, resulting in 58 datasets being analysed (see Supplementary Materials 3 for analyses). In the worsening environment, the same outlier was removed and data was missing for four people, resulting in 54 datasets being analysed (see Supplementary Materials 1 for full details of how data was processed).

Results

There was no effect of environment, with investments made in the improving environment (M 4.3, SD 1.7, CI[3.8,4.8]) being no different to those made in the worsening environment (M 4.3, SD 1.8, CI[3.8,4.8]), $f(1,50) = 0.05$, $p = .830$, $np^2 < .01$. As expected, however, the helpful face that provided 75% valid gaze-cues (M £4.7, SD £1.9, CI[4.2,5.3]) received overall higher investment amounts than the unhelpful face that only offered 25% valid gaze-cues (mean £3.9, SD £1.8, CI[3.4,4.3]), $f(1,50) = 7.53$, $p = .008$, $np^2 = .13$. Although numerically it does look as though the differences between the faces was more marked in the improving condition (**Figure 5**), perhaps due to a lack of power, the interaction was smaller than we could reliably detect $f(1,50) = 2.68$, $p = .108$, $np^2 = .05$.

In order to assess whether these potential differences are compelling enough to test in future work, we performed exploratory comparisons of simple differences between how the helpful and unhelpful faces were perceived within each environment (**Table 7**). For the

Figure 5: Investment decisions for Experiment 4 for the consistent helpful (75% valid) and stable unhelpful (25% valid) faces across the two dynamic environments; improving (left) and worsening (right). Bars represent 95% CI.

improving environment, the helpful face was trusted more than the unhelpful face, whereas there was no meaningful difference in investments made to the two faces in the worsening environment. Planned paired t-tests were used to see whether the individual faces were judged differently in the two environments, but there were no clear effects present. Both the unhelpful and the helpful faces received similar amounts, irrespective of environment.

Contrast	Mean difference	95% CI Lower	95% CI Upper	t	dz	p
Within Environments						
25-25 Improving > 75-75 Improving	-1.31	-2.15	-0.48	3.15	.44	.003
25-25 Worsening > 75-75 Worsening	-0.45	-1.28	0.38	1.09	.15	.282
Between Environments						
25-25 Improving > 25-25 Worsening	-0.37	-1.16	0.41	0.96	.13	.344
75-75 Improving > 75-75 Worsening	0.49	-0.25	1.23	1.33	.19	.189

Table 7: Pairwise comparisons between conditions both within and between environments for Experiment 4.

Discussion

Our results show small differences in how people are judged based on changing group dynamics, but these differences were not entirely in line with our predictions. We predicted contrast effects, whereby the changing valence of the wider group environment would exaggerate the effect of opposingly valenced target individuals' behaviour on impression formation. Our results provide some evidence in support of this for the unhelpful person (25% valid) but show the opposite result for the helpful person (75% valid). Indeed, trust towards the helpful person appears to have been boosted by consistently being helpful in an improving environment. Our results give some, though certainly not unequivocal, support that group dynamics may operate on impression formation differently for individuals who are perceived positively vs. negatively. Perhaps unsurprisingly, group effects on impression formation, if truly present, are smaller than the effects due to differences in an individual's own behaviour. As such, our study was not properly powered to support (or deny) these smaller effects. Thus, at present, it is unclear what might be driving these possible group dynamic effects, and we encourage future research to probe the question further.

Experiment 5

Introduction

Experiment 5 used the converse approach to that of Experiment 4. Here, it was the groups' behaviour that remained consistent and two target individuals who changed their behaviour to help more or less over time. As we designed and completed Experiment 5 in parallel with Experiment 4, we had similar general predictions. We expected contrast effects to manifest in two distinct ways. First, we expected both individuals to be perceived as less trustworthy in a helpful than an unhelpful environment, due to the fact that the overall unhelpful face, despite increasing how much it helps over time, is still helping less than other groups members and a face that helps less over time, despite helping a lot overall, is still

deviating from the group. Further, we expected increased helpfulness over time to have a larger impact in the context of the unhelpful group than the helpful group (i.e., the face that increases in helpfulness will be judged more positively in the unhelpful environment).

Likewise, we expected decreased helpfulness to have a larger impact in a generally helpful environment compared to an unhelpful environment.

Method

Participants

In return for course credit, a total of 57 people (female 40, missing 7, mean age 20.8, SD 3.4) took part in this study, although not all completed both sessions. Again, demographic information was collected in the second session of the study, meaning that the information is not available for those participants who did not attend both parts.

Design

Two faces varied in the proportion of valid cues they displayed overall and the order in which they provided their helping behaviour. The first face started by giving only valid gaze cues but then decreased to only provide 50% valid cues, resulting in 75% average helping behaviour. Conversely, the second face started by providing no valid gaze cues before increasing to offer 50% valid cues, resulting in 25% average helping (**Table 8**). These same behavioural profiles were used in Experiments 1 and 2, where we found no difference between how they were perceived. Any differences in how they are perceived in Experiment 5, then, are most likely to be due to the different social environments in which they are seen.

Each face was encountered among two contrasting social environments, one where faces consistently displayed mostly valid cues (helpful environment) and another where faces consistently provided mostly invalid cues (unhelpful environment) (**Table 8**). Each block, therefore, had an unequal number of valid and invalid trials. Given the 'stable' nature of each Environment, this difference was constant across blocks, with more valid trials being

Validity	Target Decreasing	Target Increasing	Helpful	Helpful*	Helpful
Initial (blocks 1&2)	100	0	50	75	100
End (block 3&4)	<u>50</u>	<u>50</u>	<u>50</u>	<u>75</u>	<u>100</u>
Average	75	25	50	75	100
			Unhelpful	Unhelpful*	Unhelpful
Initial (blocks 1&2)	100	0	50	25	0
End (block 3&4)	<u>50</u>	<u>50</u>	<u>50</u>	<u>25</u>	<u>0</u>
Average	75	25	50	25	0

Table 8: Conditions for the consistently helpful environment (top) and the consistently unhelpful environment (bottom) in Experiment 5. * this condition was portrayed by two faces.

present than invalid in the mostly helpful environment (2:1) and fewer being displayed in the unhelpful environment (1:2).

Data analysis

A 2 (person: increase, decrease) x 2 (environment: helpful, unhelpful) repeated measures Anova was used together with planned paired comparisons to estimate key differences between individual faces across environments.

Although 57 people attended session one, only 52 completed the unhelpful environment and 55 the helpful environment, with 51 completing both after data for one participant was removed as they accidentally completed the same version twice. There were no response outliers. One participant was removed as a gaze-cue outlier due to poor performance in the helpful environment, resulting in 50 being carried forward into the social judgement analyses (see Supplementary Materials 4 for Niceness analyses).

After the removal of one gaze-cue outlier from the helpful environment, gaze-cue analyses included 54 datasets for the helpful environment and 52 datasets for the unhelpful environment (see Supplementary Materials 3).

Results

In line with predictions, faces were generally trusted more in the unhelpful (M £5.5, SD 1.9, CI[4.0,5.1]) than the helpful (M £3.79, SD £1.56, CI[3.3,4.2]) environment $f(1,49) = 8.77, p = .005, \eta^2 = .15$. There was also an effect of person, with the increasing (0-50), yet least helpful, face (M £3.89, SD 1.6, CI[3.4,4.4]) receiving overall lower investments than the decreasing (100-50), yet most helpful, face (M £4.44, SD 1.8, CI[3.9,5.0]), $f(4.52, p = .039, \eta^2 = .08$. This appeared to be driven by the least helpful face being trusted less in the helpful environment (see **Figure 6**). There was no interaction present $f(1,49) = 0.50, p = .482, \eta^2 = .01$, though this may be due to a lack of power to detect small differences.

Exploratory comparisons (**Table 9**) found that in the unhelpful environment, there was no substantial difference between the two faces. In the helpful environment, however, the increasing (overall unhelpful) face was trusted slightly less than the decreasing (overall helpful) face. Planned paired sample t-tests explored environmental differences for each face individually (**Table 9**). The decreasing face was trusted similarly in both the unhelpful environment, where it helped the most, and the helpful environment, where it did not, though the distribution of investments does seem to differ between environments (**Figure 6**). In contrast, the face that increased was viewed differently depending on the environment, as it was trusted more in the unhelpful environment.

Discussion

The two key profiles here (100% dropping to 50% vs 0% increasing to 50%) were also included in experiments 1 and 2, where they were perceived similarly. Here, they are perceived differently, depending on the group environment. Taken together, they were trusted more in the unhelpful than helpful environment, and when group members' behaviour was mostly negative, both faces were again perceived similarly. However, when placed in a helpful environment of mostly valid cues, the face that (though increasing over time) provided the least help overall was trusted less than the decreasing (yet overall helpful) face.

Figure 6: Violin plots showing trust decisions for each target condition in Experiment 5 for the helpful (left) and unhelpful (right) social environments. Bars represent 95% CI.

This makes intuitive sense, given that in a helpful environment, helping is commonplace and so a face that deviates from the “common behaviour” may stand out more. It is unclear, however, why a similar contrast effect does not occur for the overall more helpful face in the unhelpful environment. Taken together, the results suggest a recency effect in the unhelpful environment and an averaging effect in the helpful environment. However, much like in Experiment 4, larger sample sizes are required to adequately probe the effects of social context.

Contrast	Mean difference	95% CI Lower	95% CI Upper	t	dz	p
Within Environments						
100-50 > 0-50 (Helpful group)	0.36	-0.38	1.10	0.98	0.14	.334
100-50 > 0-50 (Unhelpful group)	0.74	-0.02	1.50	1.96	0.28	.055
Between Environments						
100-50 Helpful > 100-50 Unhelpful	-0.56	-1.30	0.18	-1.53	0.22	.133
0-50 Helpful > 0-50 Unhelpful	-0.94	-1.69	-0.19	-2.53	0.36	.015

Table 9: Pairwise comparisons, both within and between environments, for the inconsistent, target conditions in Experiment 5.

General Discussion

We demonstrate several ways in which incidental impression formation is sensitive to eye-gaze cues that change over time. Our results show that despite being task-irrelevant, people track and integrate inconsistent and varied eye-gaze behaviour over time, which they use to form nuanced judgements of people's trustworthiness. In addition, we demonstrate that in the context of non-verbal social cues, impression formation and updating is not dominated by a single cognitive process but is a complex product of many cognitive mechanisms that integrate average valence over time, the direction of behaviour changes, the recency of observed behaviour, and the group context in which the behaviour is observed. Our results, therefore, not only demonstrate just how tuned humans are to changes in non-verbal social information, but that our social assessments change flexibly in response to many different factors, including group context.

The results extend prior studies that used a similar gaze-cueing paradigm to manipulate trust (Bayliss & Tipper, 2006; Manssuer et al., 2015; Rogers et al., 2014; Strachan et al., 2017; Strachan & Tipper, 2017). Compared to such prior work, where people experienced only wholly valid and wholly invalid gaze-cues, we show that people are also capable of both tracking inconsistent and less categorical depictions of social behaviour, and of tracking *changes* in that behaviour over time. The most recent behaviour was generally the best predictor of subsequent impressions. However, we also show that the relative importance of any particular factor may be more or less important depending on the context and circumstances. Thus, among the factors that we manipulated, no one factor was dominant across all five experiments. Instead, a more complex picture emerged compared to prior research using the gaze-cueing paradigm.

Such added complexity was particularly apparent in two experimental contexts. First, in Experiments 4 and 5, we found some evidence that group context influences impression formation differently for individuals whose behaviour, in isolation, would be perceived as positive as opposed to individuals who would be perceived negatively. This promising initial

finding demonstrates that how we perceive an individual can be affected by behaviour of the wider group. Similarly, Lammers and colleagues (2018) found that the 'valence homogeneity' of others in the learning environment influences how an individual's behaviour is perceived. Second, in Experiments 1 and 3, faces that started out as helpful but became more deceitful over time were trusted as little as the individual that was consistently unhelpful. Thus, while new positive information does not over-write past bad behaviour, new negative information can almost entirely expunge previous positive behaviour. The observation that negative behaviour changes may have a larger effect than positive changes is consistent with previous research, which showed that faces providing only invalid cues were trusted less, but were remembered better, than valid cue providers (Bayliss et al., 2009).

Our results can also be contextualised by considering the wider literature on impression formation that extends beyond gaze-cueing studies. Both primacy and recency effects have been documented in prior impression formation research exploring order effects. Such conflicting findings may be at least partially explained by task instructions (Forgas, 2011; Hogarth & Einhorn, 1992) and differences in the weight of negative and positive information (Baumeister et al., 2001; Fiske, 1980), where negative information received after positive information updates impressions more than positive information received after negative. Indeed, this asymmetry is the case in studies using moral statements, where impressions are asymmetrically updated depending on the direction of the valence change (Mende-Siedlecki, Baron, et al., 2013; Reeder & Coover, 1986). Such studies have typically, though not always, used written descriptions of behaviour that are clearly trait-diagnostic. Here, when non-verbal eye-gaze cues are used as social signals, we find that the most recent behaviour generally has more influence than the initial behaviour on final impressions, though positive-to-negative changes are still more powerful than negative-to-positive changes. One possible practical benefit of using non-verbal and less trait-

diagnostic cues is that they may be more easily integrated over time and therefore enable both more nuanced and more gradual changes in social behaviour to be studied.

A complementary interpretation for the dominance of recency over primacy in the context of eye-gaze cues is that recent information may be particularly powerful when the social cues are neither the focus of attention nor strongly morally valenced. In the “real world”, changes in eye-gaze cues are likely to be more indicative of a current and transient mood, which is more predictive of current behaviour towards us than enduring pro- or anti-social traits (Frith & Frith, 2006). Consequently, the most recent eye-gaze behaviour may dominate because of its superior functional relevance in this context. An alternative possibility for stronger effects of recency is that transient eye-gaze cues may be less readily encoded and/or tracked explicit and highly-valenced cues, such as written descriptions of behaviour. However, there are reasons to believe that rapid decay of incidental impressions from eye-gaze cues is not the sole driver of these results. Recent studies have found that the eye-gaze cueing effects on impression formation survived for up to a week, though the strength of the effect faded (Strachan et al., 2020; Strachan & Tipper, 2017).

Limitations and Future Directions

Some of the smaller effects that we observed need to be replicated in a future study that is properly powered to their size. However, we also expect that effects that are small in the context of a lab-based experiment could have potentially larger cumulative effects over time in real social interactions (Funder & Ozer, 2019). Of course, this needs to be tested empirically in future work, but we consider this to be particularly possible with incidental social cues as they are encountered many times a day, and our data strongly suggests that people are not only exquisitely sensitive to such cues, but also capable of tracking changes in such cues across time.

It is also important to note that participants did not perceive the “neutral”, non-predictive (50-50 valid/invalid), behaviour as neutral. Instead, they generally judged such

behaviour almost as negatively as they did individuals who were providing invalid cues a majority of the time. While we calculated the gazer's behaviour to be mathematically neutral, in reality, non-predictive behaviour is wholly unhelpful, as it can never be relied upon to make predictions. In this way, it makes sense that people would form negative opinions about such individuals. We acknowledge, therefore, that there may be no "neat" way to produce truly neutral behaviour when using this sort of task.

A further consideration is that we did not use an incentive compatible design for our trust measurement (Mailath, 1987). In other words, participants did not actually receive real monetary rewards. As a consequence, investment decisions might have been different if they actually stood to lose or gain money. While we accept that absolute investment amounts may have been different had we included a monetary incentive, we believe that the relationships found between conditions are still valid. Indeed, in our pilot work (see Supplemental Materials 2), we found effects of a similar size to those found in work that did use an incentive compatible design (Rogers et al., 2014) and we consistently found large differences between valid and invalid faces in the expected direction and in spite of no risk or reward.

Finally, while the design used here is good for providing social behaviour covertly, it is not well-suited for teasing apart what kind of cognitive mechanisms underpin trait impressions from incidental eye-gaze cues, nor what element of the behaviour is driving incidental impression processes. One possibility is that behaviour directly drives impression formation. For example, invalid cues are perceived as deceitful, leading to the inference that the person making such cues is untrustworthy. Alternatively, the consequences of cue behaviour for participants could associatively drive impressions. For example, invalid cues slow participants' responses, so faces that provide invalid cues could be associated with a 'bad' outcome, or the task being more difficult (Behrens et al, 2008). In this scenario, it is our outcome that drives impressions rather than a judgement about the meaning of the

behaviour itself. In our paradigm, these two possibilities are confounded. This distinction should be considered in future work.

Conclusion

Our results demonstrate that impression formation and updating is both complex and flexible, sensitive to how eye-gaze cues change over time in a variety of ways, as well as to the group context in which such cues are observed. Our results strongly suggest that we need a more comprehensive theory for how incidentally encountered social cues are integrated with contextual cues and more explicit social knowledge to impact on impression formation across time. As such, we feel this work reflects a compelling starting point for future research.

Acknowledgements

This research was performed as part of an all-Wales Economic and Social Research Council (ESRC) Doctoral Training Centre PhD studentship (awarded to RR and KK, PhD student: RN).

References

- Anderson, N. H. (1965). Averaging versus adding as a stimulus-combination rule in impression formation. In *Journal of Experimental Psychology* (Vol. 70, Issue 4, pp. 394–400). American Psychological Association. <https://doi.org/10.1037/h0022280>
- Argyle, M., & Cook, M. (1976). Gaze and mutual gaze. In *Gaze and mutual gaze*. Cambridge U Press.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is Stronger than Good. *Review of General Psychology*, 5(4), 323–370. <https://doi.org/10.1037/1089-2680.5.4.323>
- Bayliss, A. P., Griffiths, D., & Tipper, S. P. (2009). Predictive gaze cues affect face evaluations: The effect of facial emotion. *European Journal of Cognitive Psychology*, 21(7), 1072–1084. <https://doi.org/10.1080/09541440802553490>
- Bayliss, A. P., & Tipper, S. P. (2006). Predictive gaze cues and personality judgments: Should eye trust you? *Psychological Science*, 17(6), 514–520. <https://doi.org/10.1111/j.1467-9280.2006.01737.x>
- Berg, J., Dickhaut, J., & McCabe, K. (1995). Berg et al 1995.pdf. In *Games and Economic Behavior* (Vol. 10, pp. 122–142). <https://doi.org/10.1006/game.1995.1027>
- Brambilla, M., Carraro, L., Castelli, L., & Sacchi, S. (2019). Changing impressions: Moral character dominates impression updating. *Journal of Experimental Social Psychology*, 82, 64–73. <https://doi.org/https://doi.org/10.1016/j.jesp.2019.01.003>
- Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T.-H., Huber, J., Johannesson, M., Kirchler, M., Nave, G., Nosek, B. A., Pfeiffer, T., Altmejd, A., Buttrick, N., Chan, T., Chen, Y., Forsell, E., Gampa, A., Heikensten, E., Hummer, L., Imai, T., ... Wu, H. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behaviour*, 2(9), 637–644.

<https://doi.org/10.1038/s41562-018-0399-z>

Campellone, T. R., & Kring, A. M. (2013). Who do you trust? The impact of facial emotion and behaviour on decision making. *Cognition and Emotion*, 27(4), 603–620.

<https://doi.org/10.1080/02699931.2012.726608>

Carlson, D. E., & Mae, L. (2003). The accidental tourist: Capturing incidental (versus intentional) impressions. In *Foundations of Social Cognition: A Festschrift in Honor of Robert S. Wyer, Jr.* (pp. 97–130). Lawrence Erlbaum Associates Publishers.

Carlston, D. E., McCall, T. C., McCarty, M. K., & Tay, L. (2015). On being judged by the company you keep: The effects of group consensus and target behavior on impressions of individual group members. *Journal of Experimental Social Psychology*, 60, 173–182.

<https://doi.org/10.1016/j.jesp.2015.06.001>

Chang, L. J., Doll, B. B., van 't Wout, M., Frank, M. J., & Sanfey, A. G. (2010a). Seeing is believing: Trustworthiness as a dynamic belief. *Cognitive Psychology*, 61(2), 87–105.

<https://doi.org/10.1016/j.cogpsych.2010.03.001>

Chang, L. J., Doll, B. B., van 't Wout, M., Frank, M. J., & Sanfey, A. G. (2010b). Seeing is believing: trustworthiness as a dynamic belief. *Cognitive Psychology*, 61(2), 87–105.

<https://doi.org/10.1016/j.cogpsych.2010.03.001>

Cline, M. G. (1956). The Influence of Social Context on the Perception of Faces. *Journal of Personality*, 25(2), 142–158. <https://doi.org/10.1111/j.1467-6494.1956.tb01294.x>

Colombatto, C., Chen, Y. C., & Scholl, B. J. (2020). Gaze deflection reveals how gaze cueing is tuned to extract the mind behind the eyes. *Proceedings of the National Academy of Sciences of the United States of America*, 117(33), 19825–19829.

<https://doi.org/10.1073/PNAS.2010841117>

Cone, J., Mann, T. C., & Ferguson, M. J. (2017). Changing our implicit minds: How, when, and why implicit evaluations can be rapidly revised. In *Advances in experimental social*

psychology. (pp. 131–199). Elsevier Academic Press.

Cumming, G. (2014). The New Statistics: Why and How. *Psychological Science*, 25(1), 7–29. <https://doi.org/10.1177/0956797613504966>

Dalmaso, M., Castelli, L., & Galfano, G. (2020). Social modulators of gaze-mediated orienting of attention: A review. *Psychonomic Bulletin and Review*, 27(5), 833–855. <https://doi.org/10.3758/s13423-020-01730-x>

Delgado, M. R., Frank, R. H., & Phelps, E. A. (2005). Perceptions of moral character modulate the neural systems of reward during the trust game. *Nature Neuroscience*, 8(11), 1611–1618. <https://doi.org/10.1038/nn1575>

Doherty, M. J. (2006). The development of mentalistic gaze understanding. *Infant and Child Development*, 15(2), 179–186. <https://doi.org/https://doi.org/10.1002/icd.434>

Driver, J., Davis, G., Ricciardelli, P., Kidd, P., Maxwell, E., & Baron-Cohen, S. (1999). Gaze perception triggers reflexive visuospatial orienting. *Visual Cognition*, 6(5), 509–540. <https://doi.org/10.1080/135062899394920>

Ecker, U. K. H., & Rodricks, A. E. (2020). Do False Allegations Persist? Retracted Misinformation Does Not Continue to Influence Explicit Person Impressions. *Journal of Applied Research in Memory and Cognition*, 9(4), 587–601. <https://doi.org/https://doi.org/10.1016/j.jarmac.2020.08.003>

Emery, N. J. (2000). The eyes have it: The neuroethology, function and evolution of social gaze. *Neuroscience and Biobehavioral Reviews*. *Neuroscience and Biobehavioral Reviews*, 24, 581–604.

Fareri, D. S., Chang, L. J., & Delgado, M. R. (2012). Effects of direct social experience on trust decisions and neural reward circuitry. *Frontiers in Neuroscience*, 6(OCT), 1–17. <https://doi.org/10.3389/fnins.2012.00148>

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical

- power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191. <https://doi.org/10.3758/BF03193146>
- Fiske, S. T. (1980). Attention and weight in person perception: The impact of negative and extreme behavior. *Journal of Personality and Social Psychology*, 38(6), 889–906. <https://doi.org/10.1037/0022-3514.38.6.889>
- Forgas, J. P. (2011). Can negative affect eliminate the power of first impressions? Affective influences on primacy and recency effects in impression formation. *Journal of Experimental Social Psychology*, 47(2), 425–429. <https://doi.org/10.1016/j.jesp.2010.11.005>
- Friesen, C. K., & Kingstone, A. (1998). The eyes have it! Reflexive orienting is triggered by nonpredictive gaze. *Psychonomic Bulletin and Review*, 5(3), 490–495. <https://doi.org/10.3758/BF03208827>
- Frischen, A., Bayliss, A. P., & Tipper, S. P. (2007). Gaze cueing of attention: visual attention, social cognition, and individual differences. *Psychological Bulletin*, 133(4), 694–724. <https://doi.org/10.1037/0033-2909.133.4.694>
- Frith, C. D., & Frith, U. (2006). How we predict what other people are going to do. *Brain Research*, 1079(1), 36–46. <https://doi.org/https://doi.org/10.1016/j.brainres.2005.12.126>
- Funder, D. C., & Ozer, D. J. (2019). Evaluating Effect Size in Psychological Research: Sense and Nonsense. *Advances in Methods and Practices in Psychological Science*, 2(2), 156–168. <https://doi.org/10.1177/2515245919847202>
- Gigerenzer, G. (2018). Statistical Rituals: The Replication Delusion and How We Got There. *Advances in Methods and Practices in Psychological Science*, 1(2), 198–218. <https://doi.org/10.1177/2515245918771329>
- Gobel, M. S., Kim, H. S., & Richardson, D. C. (2015). The dual function of social gaze. *Cognition*, 136, 359–364. <https://doi.org/https://doi.org/10.1016/j.cognition.2014.11.040>

- Hamilton, A. F. de C. (2016). Gazing at me: the importance of social meaning in understanding direct-gaze cues. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1686), 20150080. <https://doi.org/10.1098/rstb.2015.0080>
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24(1), 1–55.
[https://doi.org/https://doi.org/10.1016/0010-0285\(92\)90002-J](https://doi.org/https://doi.org/10.1016/0010-0285(92)90002-J)
- Lammers, J., Gast, A., Unkelbach, C., & Galinsky, A. D. (2018). Moral Character Impression Formation Depends on the Valence Homogeneity of the Context. *Social Psychological and Personality Science*, 9(5), 576–585. <https://doi.org/10.1177/1948550617714585>
- Langton, S. R. H., & Bruce, V. (1999). Reflexive visual orienting in response to the social attention of others. *Visual Cognition*, 6(5), 541–567.
<https://doi.org/10.1080/135062899394939>
- Lee, V. K., & Harris, L. T. (2013). How social cognition can inform social decision making. *Frontiers in Neuroscience*, 7, 259. <https://doi.org/10.3389/fnins.2013.00259>
- Lundqvist, D., Flykt, A., & Öhman, A. (1998). The Karolinska directed emotional faces (KDEF). CD ROM from Department of Clinical Neuroscience, Psychology Section, Karolinska Institutet, 91(630), 2.
- Mailath, G. (1987). Incentive Compatibility in Signaling Games with a Continuum of Types. *Econometrica*, 55(6), 1349–1365.
- Mann, T. C., & Ferguson, M. J. (2015). Can we undo our first impressions? The role of reinterpretation in reversing implicit evaluations. In *Journal of Personality and Social Psychology* (Vol. 108, Issue 6, pp. 823–849). American Psychological Association.
<https://doi.org/10.1037/pspa0000021>
- Manssuer, L. R., Roberts, M. V., & Tipper, S. P. (2015). The late positive potential indexes a role for emotion during learning of trust from eye-gaze cues. In *Social Neuroscience*

(Vol. 10, Issue 6, pp. 635–650). Taylor & Francis.

<https://doi.org/10.1080/17470919.2015.1017114>

Maurer, C., Chambon, V., Bourgeois-Gironde, S., Leboyer, M., & Zalla, T. (2018). The influence of prior reputation and reciprocity on dynamic trust-building in adults with and without autism spectrum disorder. *Cognition*, 172, 1–10.

<https://doi.org/https://doi.org/10.1016/j.cognition.2017.11.007>

Mende-Siedlecki, P. (2018). Changing our minds: the neural bases of dynamic impression updating. *Current Opinion in Psychology*, 24, 72–76.

<https://doi.org/10.1016/j.copsyc.2018.08.007>

Mende-Siedlecki, P., Baron, S. G., & Todorov, A. (2013). Diagnostic Value Underlies Asymmetric Updating of Impressions in the Morality and Ability Domains. *The Journal of Neuroscience*, 33(50), 19406 LP – 19415. <https://doi.org/10.1523/JNEUROSCI.2334-13.2013>

Mende-Siedlecki, P., Cai, Y., & Todorov, A. (2013). The neural dynamics of updating person impressions. *Social Cognitive and Affective Neuroscience*, 8(6), 623–631.

<https://doi.org/10.1093/scan/nss040>

Mende-Siedlecki, P., & Todorov, A. (2016). Neural dissociations between meaningful and mere inconsistency in impression updating. *Social Cognitive and Affective Neuroscience*, 11(9), 1489–1500. <https://doi.org/10.1093/scan/nsw058>

Newman, L. S., & Uleman, J. S. (1990). Assimilation and Contrast Effects in Spontaneous Trait Inference. *Personality and Social Psychology Bulletin*, 16(2), 224–240.

<https://doi.org/10.1177/0146167290162004>

Open, S. C. (2015). Psychology. Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716.

Posner, M. I. (1980). Orienting of attention. *Journal of Experimental Psychology*, 32(July

1979), 3–25.

Reeder, G. D., & Covert, M. D. (1986). Revising an Impression of Morality. *Social Cognition*, 4(1), 1–17. <https://doi.org/10.1521/soco.1986.4.1.1>

Rezlescu, C., Duchaine, B., Olivola, C. Y., & Chater, N. (2012). Unfakeable Facial Configurations Affect Strategic Choices in Trust Games with or without Information about Past Behavior. *PLOS ONE*, 7(3), e34293.

Ricciardelli, P., Bricolo, E., Aglioti, S. M., & Chelazzi, L. (2002). My eyes want to look where your eyes are looking: Exploring the tendency to imitate another individual's gaze. *NeuroReport*, 13(17).

Ricciardelli, P., Carcagno, S., Vallar, G., & Bricolo, E. (2013). Is gaze following purely reflexive or goal-directed instead? Revisiting the automaticity of orienting attention by gaze cues. *Experimental Brain Research*, 224(1), 93–106.
<https://doi.org/10.1007/s00221-012-3291-5>

Rogers, R. D., Bayliss, A. P., Szepietowska, A., Dale, L., Reeder, L., Pizzamiglio, G., Czarna, K., Wakeley, J., Cowen, P. J., & Tipper, S. P. (2014). I want to help you, but I am not sure why: Gaze-cuing induces altruistic giving. *Journal of Experimental Psychology: General*, 143(2), 763–777. <https://doi.org/10.1037/a0033677>

Shen, X., Mann, T. C., & Ferguson, M. J. (2020). Beware a dishonest face?: Updating face-based implicit impressions using diagnostic behavioral information. *Journal of Experimental Social Psychology*, 86(July 2019), 103888.
<https://doi.org/10.1016/j.jesp.2019.103888>

Siegel, J. Z., Mathys, C., Rutledge, R. B., & Crockett, M. J. (2018). Beliefs about bad people are volatile. *Nature Human Behaviour*, 2(10), 750–756. <https://doi.org/10.1038/s41562-018-0425-1>

Simpson, D. D., & Ostrom, T. M. (1976). Contrast effects in impression formation. *Journal of*

Personality and Social Psychology, 34(4), 625–629. <https://doi.org/10.1037/0022-3514.34.4.625>

Skowronski, J. J., & Carlston, D. E. (1989). Negativity and extremity biases in impression formation: A review of explanations. *Psychological Bulletin*, 105(1), 131–142. <https://doi.org/10.1037/0033-2909.105.1.131>

Strachan, J., & Tipper, S. (2015). Using scalar ratings to track changes in apparent trustworthiness induced by helpful and misleading gaze cues. *Journal of Vision*, 15(12), 1216. <https://doi.org/10.1167/15.12.1216>

Strachan, J. W. A., Guttesen, A. Á. V., Smith, A. K., Gaskell, M. G., Tipper, S. P., & Cairney, S. A. (2020). Investigating the formation and consolidation of incidentally learned trust. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 46(4), 684–698. <https://doi.org/10.1037/xlm0000752>

Strachan, J. W. A., Kirkham, A. J., Manssuer, L. R., Over, H., & Tipper, S. P. (2017). Incidental learning of trust from eye-gaze: Effects of race and facial trustworthiness. *Visual Cognition*, 25(7–8), 802–814. <https://doi.org/10.1080/13506285.2017.1338321>

Strachan, J. W. A., & Tipper, S. P. (2017). Examining the durability of incidentally learned trust from gaze cues. *Quarterly Journal of Experimental Psychology*, 70(10), 2060–2075. <https://doi.org/10.1080/17470218.2016.1220609>

Uleman, J. S., Saribay, S. A., & Gonzalez, C. M. (2008). Spontaneous inferences, implicit impressions, and implicit theories. *Annual Review of Psychology*, 59, 329–360. <https://doi.org/10.1146/annurev.psych.59.103006.093707>

Wyer, N. A. (2010). You Never Get a Second Chance to Make a First (Implicit) Impression: The Role of Elaboration in the Formation and Revision of Implicit Impressions. *Social Cognition*, 28(1), 1–19. <https://doi.org/10.1521/soco.2010.28.1.1>

Zarolia, P., Weisbuch, M., & McRae, K. (2017). Influence of indirect information on

interpersonal trust despite direct information. *Journal of Personality and Social Psychology*, 112(1), 39–57. <https://doi.org/10.1037/pspi0000074>

Zwaan, R. A., Etz, A., Lucas, R. E., & Donnellan, M. B. (2018). Making replication mainstream. *Behavioral and Brain Sciences*, 41, e120. <https://doi.org/DOI:10.1017/S0140525X17001972>