Article:

Mole, CD orcid.org/0000-0002-1463-6419 and Wilkie, RM orcid.org/0000-0003-4299-7171 (2017) Looking forward to safer HGVs: the impact of mirrors on driver reaction times. Accident Analysis and Prevention, 107. pp. 173-185. ISSN 0001-4575

https://doi.org/10.1016/j.aap.2017.07.027

Mole, C. D. & Wilkie, R. M. (in press). Looking forward to safer HGVs: the impact of mirrors on driver reaction times. *Accident Analysis & Prevention. Author's Version of the Accepted Manuscript*

Looking forward to safer HGVs: the impact of mirrors on driver reaction times

Mole, C. D. & Wilkie, R. M.*

School of Psychology, University of Leeds.

*r.m.wilkie@leeds.ac.uk

Highlights

- Drivers react slower to objects viewed in mirrors compared to the front windscreen
- Drivers react slower to objects viewed in mirrors that are further from the front
- Additional cognitive load slows reaction times further for both mirrors and windows
- Cab designs which reduce the need for monitoring multiple mirrors are recommended.

Article Accepted: 28/07/17

Keywords: Driving Simulator; Vehicle Design; Heavy Goods Vehicles; Reaction Times; Cognitive Load

Abstract

Heavy Goods Vehicle (HGV) collisions are responsible for a disproportionate number of urban vulnerable road user casualties (VRU – cyclists and pedestrians). Blind-spots to the front and side of HGVs can make it difficult (sometimes impossible) to detect close proximity VRUs and may be the cause of some collisions. The current solution to this problem is to provide additional mirrors that can allow the driver to see into the blind-spots. However, keeping track of many mirrors requires frequent off-road glances which can be difficult to execute during demanding driving situations. One suggestion is that driving safety could be improved by redesigning cabs in order to reduce/remove blind-spot regions, with the aim of reducing the need for mirrors, and increasing detection rates (and thereby reducing collisions). To examine whether mirrors delay driver responses we created a series of simulated driving tasks and tested regular car drivers and expert HGV drivers. First we measured baseline reaction times to objects appearing when not driving ('Parked'). Participants then repeated the task whilst driving through a simulated town (primary driving tasks were steering, braking, and following directional signs): driving slowed reaction times to objects visible in mirrors but not to objects visible through the front windscreen. In a second experiment cognitive load was increased, this slowed RTs overall but did not alter the pattern of responses across windows and mirrors. Crucially, we demonstrate that the distribution of mirror RTs can be captured simply by the mirror's spatial position (eccentricity). These findings provide robust evidence that drivers are slower reacting to objects only visible in eccentric mirrors compared to direct viewing through the front windscreen.

Introduction

Across the world, heavy goods vehicles (HGV)[1] collisions are consistently over-represented in vulnerable road user casualties (cyclists and pedestrians). Manson et al., (2012) showed that in London 20% of cyclist casualties requiring full trauma-team activation (2004-2009) were cyclist-HGV collisions, despite HGVs travelling only 4.5% of the total distance driven on London's roads over the same time period (Traffic Analysis Centre, 2012). Other complementary studies looking at collision statistics in different countries have concluded that colliding with an HGV has particularly stark consequences for VRUs (i.e. high chance of fatality): in UK (Gilbert and McCarthy, 1994), Czech Republic (Bíl et al., 2016), USA (Ackery et al., 2012), and Germany (Niewoehner and Berg, 2005). Whilst these collisions are a huge problem on a global scale, it seems unlikely that HGVs and VRUs will stop road-sharing any time soon, since urban freight transport is critical to city economies, and cycling in major cities is becoming increasingly popular (Transport for London, 2015). Consequently, increasing the safety for cyclists sharing roads with HGVs is a topic of growing international focus (Davis and White, 2015; Pattinson and Thompson, 2014; T&E, 2016).

One well documented explanation for the over-representation of HGVs in VRU collisions lies in the limited field of view of the driver. The height of an HGV driver cab can be around 2.5m off the ground, uniquely high compared to most other motor vehicles (Niewoehner and Berg, 2005; Summerskill et al., 2015). In theory a high vantage point could offer an HGV driver an unrivalled view on the scene, enabling the driver to see across other traffic and over various obstacles. However, when a driver's vertical field of view is clipped by the bottom of the windscreen or window then a high eye-height becomes problematic, leading to large regions of non-visible space close to the vehicle (Figure 1). It has been suggested that these regions of non-visible space ('blind-spots') are critical to explaining many HGV-VRU collisions in urban environments (Cheng et al., 2016; Keigan et al., 2009; Niewoehner and Berg, 2005; Summerskill and Marshall, 2014).

[1] The European classification of HGVs is any vehicle with a total weight above 3.5 tonnes. This manuscript is particularly relevant to high-cab commercial freight vehicles which require compulsory blind-spot mirrors as per Directive 2003/97/EC.

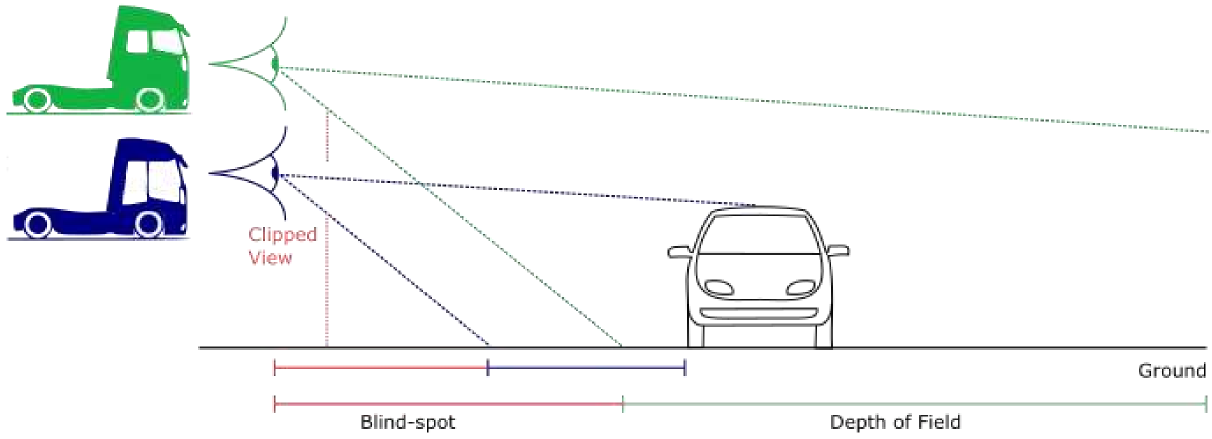

Figure 1. Raising the eye-height of the driver increases the depth of field of view, but when the field of view becomes clipped by a windscreen or window large blind-spots close to the vehicle are formed.

The current solution to improving the driver's view of the space around their vehicle is to add 'blind spot mirrors' to the vehicle so that the driver can see into regions to the front or side of the vehicle which would otherwise not be directly visible. These mirrors, if set up correctly (which is not always the case – Delmonte et al., 2012), reduce but do not eliminate blind-spots (Summerskill et al., 2015).

When assessing the usefulness of blind-spot mirrors there are a number of perceptual-motor implications to consider. For the driver making a complete scan of all mirrors requires multiple oculomotor movements (there are commonly at least six mirrors around an HGV cab) in addition to monitoring the scene directly via windows, as well as any in-vehicle information systems (e.g. GPS) in the cab (henceforth we will refer to all these possible 'windows' of information as 'viewports'). Some eye movements, such as saccadic eye movements to the same depth plane, have small latencies (time until initiation of movement) and can be executed extremely quickly (Leigh and Zee, 2006; Yang et al., 2002). However, typically a driver will be making oculomotor movements to targets that vary in location, which may require vergence eye-movements to keep the stimuli focussed and a combination of eye-head (gaze) movements to enable an eccentric (i.e. peripheral) target to be fixated. Combined saccadic-vergence eye movements, for example switching between far and near focus to look from the far road to an in-cab visual display unit, are slower than pure saccades (Leigh and Zee, 2006; Yang et al., 2002). Similarly, gaze latencies increase with eccentricity of the target (Goldring et al., 1996), so executing a head turn to look at mirrors in the periphery (such as the Class V side blind-spot mirrors) may be slow.

Consequently, a complete scan of all viewports available to the driver may be time consuming. It is difficult to precisely estimate how long a complete scan would take. In an on-road driving task in cars, Sodhi, Reimer, & Llamazares (2002) found that execution time from fixating the road in front of the vehicle to fixating some in-car object (a rear-view mirror, radio, or odometer) was approximately .33s (ranging from .23s to .47s). Similar measures of execution time have not been taken for individuals driving HGVs, and it MIGHT be that HGV mirrors are somewhat quicker or slower to fixate than a car rear-view mirror (HGV mirrors are not as close to the driver so require less vergence changes, but they are also more eccentric so require greater gaze shifts and head turns). Even if one presumes that HGV drivers can respond with the quickest execution times reported in Sodhi et al. (2002) and then adds the expected latencies (around 200ms, Leigh & Zee, 2006) to the mirror dwell time (~250ms), the driver will be taking approximately 4s to retrieve information from all six mirrors (see also Woolsgrove, 2014, for a similar estimated time). So whilst use of the mirrors may reduce the *spatial blind-spots,* the time it takes to sample from them all creates a *temporal blind-spot*, whereby the visual information viewed in the first mirror may be valid no longer (i.e. a VRU may have appeared, unnoticed, in that mirror). Given that the transport safety literature recommends an upper limit of 1.6s for off-road glances (Horrey and Wickens, 2007; Wierwille, 1993), this minimum scan time seems hugely problematic from a road safety perspective.

The temporal costs of sampling from multiple viewports have been highlighted, but there may be also visual processing costs associated with viewing objects in a mirror rather than through a window. In order to increase the spatial zone viewable by the driver blind-spot mirrors usually have a wide-angle view, which means that viewed objects are optically compressed. People respond more slowly to objects that are smaller and appear distorted compared to when they are viewed directly; indeed the smaller the object the slower the reaction time (Osaka, 1976). Even in distortion-free mirrors, participants are slower at responding to changes of objects compared to when the images are not reflected (Sareen et al., 2015). One reason for reduced reaction times may simply be the indirect mapping between the visuospatial information presented in the image, and the location of the image in the world. For example it has been shown that viewing objects indirectly (i.e. via VDUs that were not distorted or mirrored), disrupts visual-motor mappings and consequently slows motor responses (White et al., 2016). It seems, then, that drivers may respond more slowly to objects if they are viewed via a mirror compared to if they are viewed via a window, and this could itself make collisions more likely in HGVs where there is huge reliance on the use of many mirrors to cope with blind-spots.

There appears to be a solid rationale for predicting that drivers will be slower to respond to objects viewed via mirrors compared to windows. However, in situations where driving complexity is low drivers may be able to compensate by performing pre-learned scanning strategies linked to particular manoeuvres. In real-world driving HGV drivers often have to cope with highly demanding situations: consider navigating a busy road intersection in a major European city. In this situation the driver needs to distribute their cognitive resources across many tasks, keeping track of unpredictable traffic and VRUs whilst also navigating the route successfully using road signs or an in-vehicle satellite navigation system in order to make a delivery on time (Delmonte et al., 2012). In addition to those spatial and temporal constraints, there are many metrics of performance which drivers are incentivised to optimise, such as fuel efficiency (Delmonte et al., 2012). All these tasks compete for cognitive resources, and if the driver reaches capacity then the ability to successfully scan mirrors and windows will be impacted (Engström et al., 2017, 2005; Victor et al., 2005). Such conditions are often referred to as 'cognitive overload'. It is predicted that under situations of high cognitive load any existing slowing of response between windows and mirrors would be exacerbated.

Under simple reaction time (RT) conditions (where the sole task of the participant is to rapidly respond to an appearing object) humans can respond very quickly indeed (200-250ms), though reactions times increase with greater task complexity, e.g. choice reaction times (Deary et al., 2011). Similar patterns are observed in driving, where much of the reaction time research has focussed on braking responses. In a review of the literature, Green (2000) concluded that under optimal conditions (i.e. participants are expecting a hazard and there is good visibility) brake RTs are around .75s (including .2s movement time from accelerator to brake). Importantly, RTs provide a useful measure of task complexity as they increase with cognitive load (Green, 2000), such as using in-car devices (Summala et al., 1998) or otherwise responding to a concurrent task (Levy et al., 2006). RTs are therefore a suitable method of assessing whether drivers are able to respond in a timely and appropriate manner to hazards viewed across different viewports.

The current manuscript sets out to examine whether the capacity of a driver to respond quickly to an appearing stimulus varies across viewports. We also probe whether the pattern of responses interact with level of detection difficulty (high or low visibility targets), or additional cognitive load. Since harder-to-spot targets need to be fixated to be seen (i.e. low contrast objects are not easily detected in the visual periphery) changing object visibility places greater demands on visual scanning of the driver. Experiment 1 examined this issue in both regular car drivers and a sub-group of expert HGV drivers. In Experiment 2, we investigated whether an additional cognitive load (analogous to monitoring an in-vehicle visual display) impaired detection of objects differentially across viewports.

Experiment 1: Detecting Objects in Mirrors

Methods

Participants

A sample of 30 University staff (who all held a UK driving license; mean age = 28.6 years, range = 21-48yrs, 20 females), and a sample of 11 HGV drivers recruited from UK commercial freight companies (who all held a heavy goods operator licenses and drove category N_2 (3.5 – 12 tonnes) and/or N_3 (>12 tonnes) commercial vehicles as part of their occupation; mean age= 45.9 years, range = 27-59yrs, all males) were recruited for Experiment 1. All participants were naïve as to the precise purpose of the experiment. The research complied with the Declaration of Helsinki All participants gave written informed consent, and the study was approved by the University of Leeds School of Psychology Ethical Committee (Ref: 16-0096).

Apparatus

A virtual environment was programmed using WorldViz Vizard 3.0 (WorldViz, Santa Barbara, CA) on a PC with Intel i7 3770 (3.40 GHz), and projected (using two EPSON EH-TW5210 projectors) onto an L-shape half-CAVE: a forward-facing and side-facing view (see Figure 2). The refresh rate was 60Hz, with reaction time data sampled at 30Hz. The front face projection subtended 1.67m x .9m (79.72° x 48.46° field of view), whilst the side projection subtended .8m x 1.1m (43.6° x 57.62). Participants sat on a height-adjustable driving seat (physical eye-height was 1.2m, but the simulated eye-height was 2m) in a matte-black viewing booth (so that the projectors were the sole source of light), and controlled steering and braking using a force-feedback wheel and pedals (Logitech G27, Logitech, Fremont, CA) with self-centring torque .

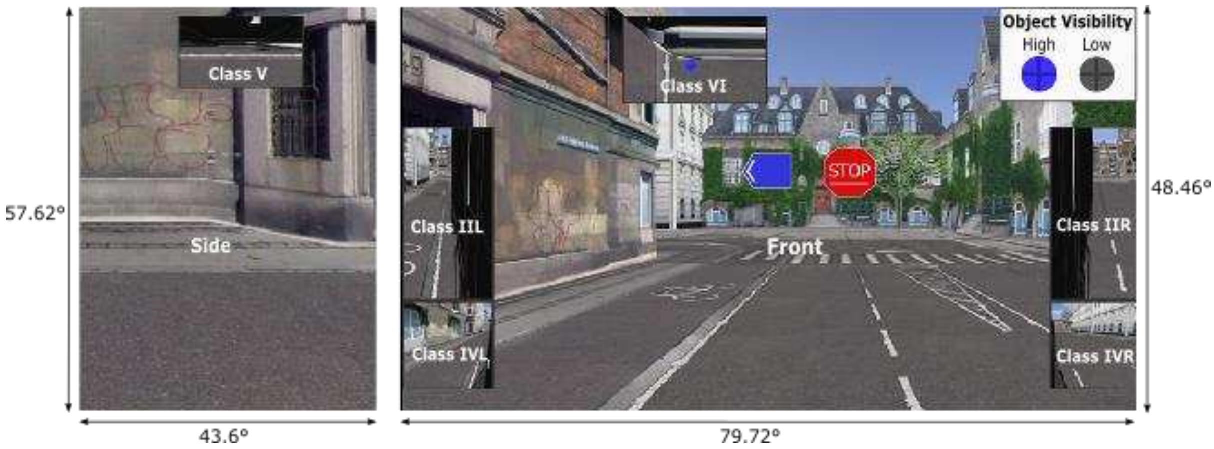

Figure 2. Simulator screenshots annotated with labels for windows and mirrors and high-visibility (Blue) and low-visibility (Grey) objects enlarged and super-imposed onto the screenshot (top-right hand corner). The high-visibility object is visible in the driver's Class VI mirror. The screenshot shows a driver approaching a zebra crossing. The 'Stop' sign instructs the driver to come to a complete stop until the 'Stop' sign flips to 'Go', whereupon they can turn move off and turn left.

Stimuli

Virtual Environment

Participants were required to drive round a virtual city, adhering to signage that told them where to go and whether to stop at a junction (Figure 2). Eight viewports were rendered simultaneously: two windows, and six mirrors (Figure 2). Rather than simulating one particular make of cab (and potentially having results that would not generalise to other models of cab) we established dimensions of viewing characteristics in the following way. The optical extent of the mirrors were approximated by taking the average of five commercially available mirrors of each class, and the ground coverage of each mirror was setup to adhere to the minimum viewing requirements set out in Directive 2003/97/EC. This means that the simulated viewports had similar viewing angle relationships (such as wide-angle Class IV mirrors) to their real-world counterparts, so that the optical extent of the scene viewed via the mirrors would been subject to similar optical scaling as real-world mirrors. This set up did not attempt to precisely match the *placement* of the mirrors in a real-world HGV (not least because mirror placement can vary from vehicle to vehicle based on individual driver preference), but our approximate placings ensured that the key characteristics relevant to rapid responses were simulated (e.g. the Class V mirror is sufficiently eccentric so that fixation requires co-ordinating eye-movements with a head turn).

Target Objects

We were interested in whether driver reaction times depended on the spatial location of the viewport. To avoid the potential confounds of reflected objects (e.g. Sareen et al., 2015) or changing object size (e.g. Osaka, 1976) also reducing reaction times, circular symmetrical targets were used that were a fixed optical size of 1.95 degrees^2, independent of the viewport (Figure 2, inset). We used high-visibility (blue) and a low-visibility (grey) objects (Figure 2, inset). Levels of visibility were piloted to be easily detected in the periphery (blue), or not (grey), with low-visibility requiring gaze fixation to be seen. The low-visibility object therefore placed greater demands on the oculomotor system, since eye-movements were needed to sample directly from all viewports.

Procedure

The participant's task was to respond to appearing objects as quickly as they could using the 'paddle buttons' behind the steering wheel (designed for use as gear shifters). A simple paddle-press was chosen to reduce the possibility of large movement times inflating reaction time (the paddles were less than 10cm away from the hand so the movement component is likely to be <222ms, Simonen, Battie, Videman & Gibbons, 1995), and so that the response mode did not interact with the other concurrent tasks (e.g. the ability to drive). The object had an equal probability of appearing in any viewport. There was an irregular time interval between responses and the appearance of the next object that was at least 2 seconds long. First participants completed the task whilst parked (not driving) with the high-visibility (HV) and low-visibility (LV) targets to obtain baseline RTs. They then repeated these experiments whilst also driving, where they were required to follow the signed route and stop at junctions if instructed (Figure 2). For the driving task the HV and LV targets were alternately interleaved to mitigate practice effects. Furthermore, since we expect slower RTs for LV targets participants always completed LV after HVs, so that a quicker RT for HV targets could not be attributed to practice effects. The total experiment time was approximately 30 minutes, split into the following blocks:

- *Parked Reaction Time: High-Visibility (6 trials x 8 Viewports = 48 trials).*

- *Parked Reaction Time: Low-Visibility (48 trials).*
- *Driving Reaction Time: High-Visibility A (3 trials x 8 viewports = 24 Trials)*
- *Driving Reaction Time: Low-Visibility A (24 Trials)*
- *Driving Reaction Time: High-Visibility B (24 Trials)*
- *Driving Reaction Time: Low-Visibility B (24 Trials)*

Analysis

The primary measure was Reaction Time in seconds from object onset to paddle button press. In order to remove unrealistically quick responses, RTs quicker than 100ms were treated as errors and removed (Luce, 1986). This removed 3 responses in the Parked task, and 3 responses in the Driving tasks. Since the target stayed on-screen until the participant responded, there were also occasional outlying trials where the participant did not notice the target for a long time (the most extreme case across all experiments was 64.2s). We wished to avoid these outlying trials disproportionally skewing the calculated means, but also wished to capture trials where there was a complete failure of detection (which would be potentially fatal for a VRU in the real-world). We opted for using a cutoff that removed the slowest 1% of RTs (within the 5% maximum cutoff suggested by Ratcliff, 1993). For the Parked condition the cutoff was 7.42s, whereas for the Driving condition it was 9.44s. This led to 43 and 42 (Parked and Driving, respectively) trials being excluded across all participants. Since the RT distribution was positively skewed we applied a log transform in order to satisfy ANOVA assumptions of normality (see Ratcliff, 1993; Whelan, 2008).

In order to assess whether there were differences between mirrors and windows, individual RT estimates were taken for front and side windows and compared against a mean taken across all mirror types. There were a number of differences across the different mirror types and so response were expected to vary. This issue probed later when looking at the individual mirror characteristics and impact on RTs (see *Further Analysis – Modelling differences between Mirrors*).

We were interested in whether the pattern of responses changed across viewports and levels of visibility, but also whether having to steer (or not) interacted with these factors. Furthermore, we wished to establish if there was a difference between Car drivers and HGV drivers, to see whether our findings would generalise to HGV drivers (who are more accustomed to using many mirrors so may have learned particular scanning strategies). This led to use of a Mixed Model ANOVA with 4 factors: Three within-subjects factors of Viewport (Front, Side, Mirror), Visibility (HV, LV), and Driving (Driving, Parked), and one between-subjects factor of Driver Type (Car, HGV).

To ensure that our findings could not simply be explained by practice or fatigue during certain conditions we ran a separate Mixed Model ANOVA on the Driving conditions with Block (A, B) as an additional factor (Within: Block, Viewport, Visibility; Between: Driver Type) on the Driving condition. There was a significant effect of Block (F (1, 39) = 10.63, p=.002, η^2 = .017) suggesting that individuals improved over time. Crucially, Block did not interact with the other factors i.e. participants improved to the same extent in all conditions. Therefore we collapsed across Blocks for the larger ANOVA described above.

Eta squared effect sizes (Bakeman, 2005) are reported, alongside Greenhouse-Geisser corrections (ε) where assumptions of sphericity were violated. Statistical tests (and later modelling) were conducted in R (R Core Team, 2016) using packages ez (Lawrence, 2016) and lme4 (Bates et al., 2015) .

Results

Reaction times when parked are shown in the boxplots in Figure 3 (left panel), with means shown as diamonds. For HV objects the grand mean RT was .82s (SE=.023s). It is clear that for LV objects RT times were slower and more widely distributed (grand mean = 1.42s, SE=.049s). RTs across viewports appeared to be quickest for objects appearing in the Front window (grand mean = .99s, SE=.054s), but it is less clear whether there are differences between objects appearing in the Side window (grand mean = 1.18s, SE=.066s) compared to the mirrors (grand mean = 1.19s, SE=.048s). There also do not appear to be systematic differences between Driver Type (HGV vs. Car).

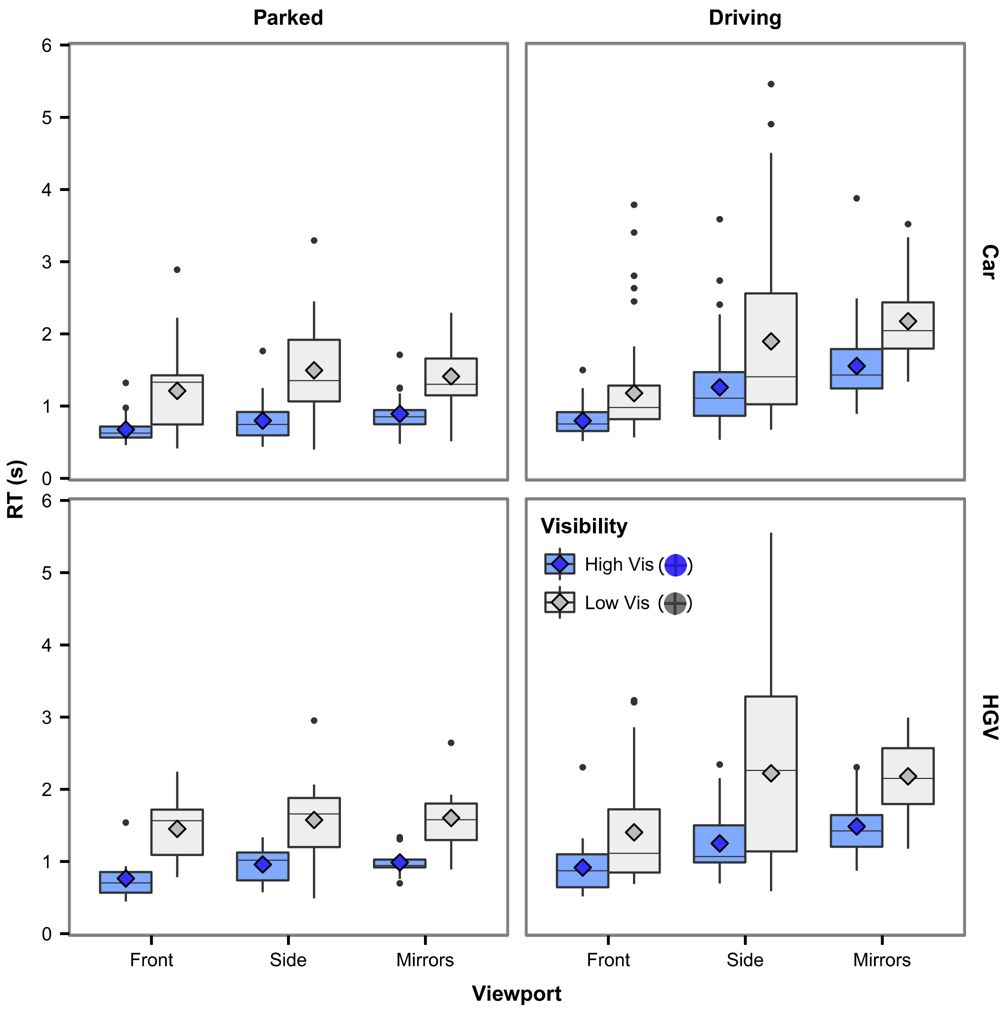

Figure 3. Boxplots showing distribution of RTs (seconds) across Viewports (x-axis) and Visibilities (Blue = high-visibility; Grey = low-visibility) for Car drivers (Top Row) and HGV drivers (Bottom Row). Left Panel = Parked. Right Panel = Driving. Diamonds represent group means.

RTs whilst Driving (collapsed across A&B blocks) are shown in Figure 3 (right panel). As with the Parked RT task, it seems clear from Figure 3 that HV objects (M=1.21s, SE=.034s) were responded to more quickly than LV objects (M=1.8s, SE=.06s). It is worth noting that RTs for HV and LV objects are approximately .4s slower when participants were required to drive as compared to when they were stationary (i.e. having to steer did not differentially affect HV or LV objects).

The ANOVA confirmed these patterns and revealed significant main effects of Driving (F (1, 39) = 70.74, p<.001, η^2 = .15), Viewport (F (2, 78) = 104.76, p<.001, η^2 = .25), and Visibility (F (1, 39) = 245.71, p<.001, η^2 = .32), but not a significant effect of Driver Type (F (1, 39) = 2.58, p=.016, η^2 = .02). The ANOVA also revealed the presence of significant interactions between Driving and Viewport (F (1.51, 58.99) = 21.43, p<.001, η^2 = .065, ε = .76) and Driving and Visibility (F (1, 39) = 8.94, p=.004, η^2 = .017).

In order to interpret the Driving x Viewport interaction the within-individual difference between Driving and Parked means are plotted in Figure 4A, across levels of Viewport (means were collapsed across visibility). Whilst in general RTs slowed when participants were required to drive, post-hoc comparisons suggest that slowing did not occur for objects appearing through the front window (seen by a difference centred on zero in Figure 4A, p=.69). RTs were significantly slower, however, for both Side (p<.001) and for Mirrors (p<.001), but the slowing effect was considerably more variable for objects presented in the Side Window.

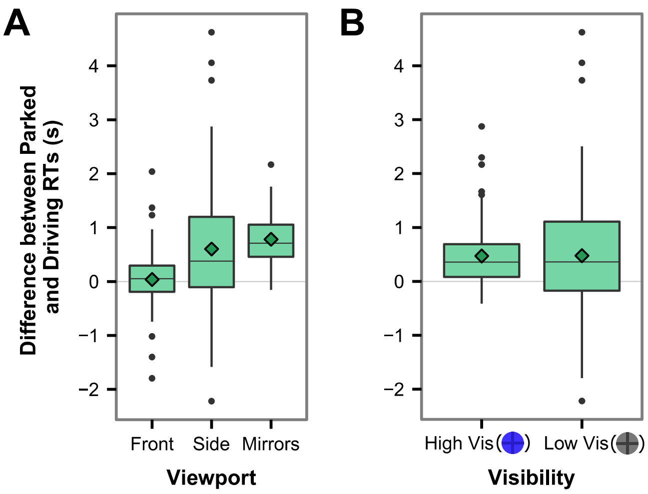

Figure 4. Investigating the interactions between Driving and Viewport (A) and Driving and Visibility (B). Boxplots show the distributions of difference between means of Parked and Driving RTs. Diamonds represent mean differences.

The Driving x Visibility interaction is plotted in Figure 4B. For both HV (p<.001) and LV objects (p=.001), driving slowed RTs to a similar extent (on average), but the interaction lies in the variability of the difference across participants. Driving had a considerably more variable effect on drivers' ability to respond to LV objects than for HV objects.

Discussion

When participants were not required to drive (Parked) they were quickest at responding to objects appearing in the Front window and slower to respond to objects appearing in the Side window or the mirrors. Driving exacerbated these differences, further slowing RTs to objects presented on the mirrors or the Side, but not slowing RTs to objects presented in the Front window. This pattern of results was consistent across both Car and HGV drivers, and also high and low visibility objects – although responses to low-visibility objects were overall slower and more variable.

The interaction between Viewport and Driving may be explained by the changing gaze requirements of the task when driving as compared to when stationary. During the Parked reaction task participants are free to adopt a wide range of visual search strategies since there are few competing demands for visual attention. Indeed, McKenzie & Harris (2015) have demonstrated that participants sample more widely from a scene when they are not required to actively steer. Since participants are able to distribute fixations equally across viewports, differences in RTs across viewports are small when not driving.

Introducing the driving (steering and navigation) task adds competing visual search demands since drivers need to sample from the scene ahead to anticipate future steering requirements (Wilkie et al., 2008; Lehtonen et al., 2013). The information necessary for the driving task (navigation and staying in lane) can be sampled through the front window (e.g. signage, lane markings), but it is less clear how information sampled through the mirrors or side window is directly relevant to the steering and navigation components of the task (though it could be argued that the side window can sometimes be useful for planning steering manoeuvres when turning at junctions). It is reasonable to assume that participants' gaze was more centralised on the front window during the driving reaction time task rather than the Parked reaction time task (cf. McKenzie & Harris, 2015). Regions in the visual field near the point of fixation are sampled with higher-acuity than regions in the periphery (Osaka, 1976), therefore the nearer an object appears to where the driver is already looking the greater the chance of rapidly detecting the stimuli – it follows that a more centralised gaze pattern would disproportionately slow reaction times to objects appearing on the eccentric viewports (side window & mirrors) but not for the front window.

Although the placement of the simulated viewports was not designed to exactly match the real-world placement of HGV drivers' own mirrors, it might have been predicted that HGV drivers would be well-practiced at monitoring many mirrors and so not exhibit differences between viewports to the same extent as Car drivers. However, despite the smaller sample of HGV drivers the Viewport x Driving interaction was consistent across both Car and HGV drivers, demonstrating that the behaviours/responses observed in regular drivers was not simply because participants were unaccustomed to the a distributed layout of many viewports.

It is striking that robust differences between viewports emerged using such a simple driving task. The simulation was not particularly 'busy' (by city driving standards, with no other road users and using just a simple navigation task). In the real-world, drivers need to cope with complex road intersections and demanding visual scanning requirements distributed across the whole scene, both outside and inside the cab (e.g. to monitor a GPS). The next experiment examined whether an additional cognitive task presented via an in-vehicle display unit selectively impairs the ability to respond to objects appearing in mirrors and windows.

Experiment 2: Additional Cognitive Load

Methods

In Experiment 2 we wished to examine whether adding a distractor task which requires both *frequent glances* and *attentional resources* increased reaction times to objects appearing in particular viewports.

Participants

A sample of 30 Car drivers were recruited through the University of Leeds staff participant database. All held a UK driving license (mean age = 29.53, range = 20-50yrs, 19 females). Unfortunately three participants withdrew from the experiment due to experiencing motion sickness. One participant did not respond to the distraction task when driving so their data was also removed (leaving N=26). All participants were naïve as to the precise purpose of the experiment.

Stimuli

The same design as Experiment 1 was used with respect to participants responding to objects in viewports. The additional Distractor Task used was a digit-pair identification task, adapted from Hines (1990). A small visual display unit was placed at the bottom of the front display (similar to having a VDU – see Figure 5). The box displayed a pair of digits which changed every two seconds. The participant's task was to check the displayed digits and respond when *both digits were odd* (e.g. "1 3" or "9 5") by clicking one of the red buttons on the front panel of the steering wheel. They did not need to respond when the digits were not both odd (e.g. "1 4" or "6 8"). Hines (1990) demonstrated that simple response times to odd-odd digit pairings are around .93s, so the participant would need to frequently check the visual display unit in order to leave themselves enough time to accurately respond (i.e. they would have to look at the digits within the first second of presentation).

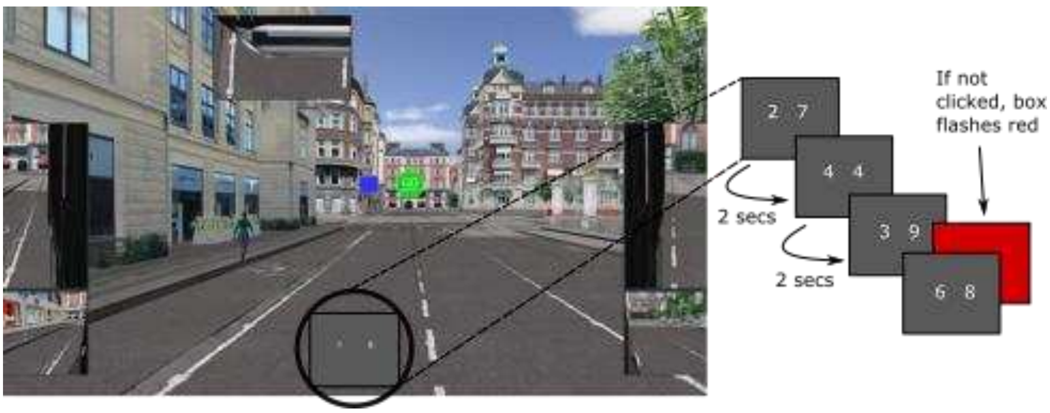

Figure 5. Number pairs changed every 2 seconds (right panel) and the box briefly flashed red if the participant made an error (clicking when pairings were odd-even or even-even, or not clicking for odd-odd pairings). The distractor task

requires the driver to look, think, and respond, so mimics driver distractions caused by an in-vehicle information system such as a phone or GPS.

If the participant missed an odd-odd pairing, or falsely responded when the digits were not both odd, then an error was recorded and the visual display unit briefly flashed red (Figure 5, right panel). The probability of an odd-odd pairing was 31.3%.

Changing the digit pair at regular intervals requires the driver to look at the visual display unit intermittently. The participant therefore has to decide when to look at the visual display unit and risk being slow in detecting an object appearing in the viewport (the other concurrent task). These tasks force drivers to engage in a trade-off between better performance at the distractor task (i.e. prioritise looking at the distractor task and responding appropriately to the odd-odd pairings) and better performance at the reaction time task (prioritising looking to the viewports and responding rapidly to targets). Since dual-task trade-offs can be mediated by task instructions (Jansen et al., 2016), the participants were asked to *prioritise the task equally with other tasks they were completing at the same time.*

Analysis

Distractor Task

To assess Distractor Task performance two measures were obtained per block: average Distraction reaction time (seconds; RT_D), and percentage correct (PC_D). When calculating RT_D only correct responses were analysed (i.e. false positives were removed). The RT_D measure was normally distributed (i.e. there was not a long tail since the maximum RT_D was 2s, the time when the numbers changed on the VDU) so the mean raw RT_D was taken (after a minimum RT cutoff of 100ms was taken to remove spurious pre-emptive responses; Luce, 1986).

Before examining whether the distractor task influenced object detection it is important to establish the extent to which participants were engaging with the distractor task. RT_D and PC_D for all trials are plotted against one another in Figure 6A. Participants responded appropriately to odd-odd pairings most of the time (PC_D mean=82.8%, SE=1.1%; RT_D mean = 1.28s, SE = .014s), and it is clear from Figure 6A that there is a strong negative correlation between RT_D and PC_D (Pearson's R = -0.53), suggesting that people who were slower to respond were also less likely to be correct. It might have been predicted that participants would try to minimise the amount of attentional resources allocated to the distraction task by adopting strategic heuristics. For example, participants could have simply responded whenever the number pair changed, resulting in quick RT_D and 100% PC_D, but also a large amount of false positives (responses when the number pairs where odd-even or even-even). However, false positives were rare – on average only 1.1 false positives (SE=.12) per block – suggesting a 'click-on-change' strategy was not adopted. Alternatively, occasional responding (guesswork) would have resulted quick RT_D (good performance) but also low PC_D (bad performance). Rather than using such simple heuristics, it appears that individuals were following instructions (to prioritise the distractor task equally with other tasks) and trying to respond both rapidly and accurately to the distraction task. Since poor performance on either measure (i.e. high RT_D or low PC_D) reflects individuals finding the distractor task difficult, we combined RT_D and PC_D into a composite measure: $Comp_D$ (given by the formula RT_D/PC_D). Since a lower PC_D score would penalise (increase) RT_D proportional to the distance of PC_D from perfect performance (100% correct), lower $Comp_D$ scores reflect better performance (see Figure 6B).

Object Reaction Time Task

For the object reaction time task (RT_O) the same pre-processing actions as Experiment 1 were taken (see *Experiment 1, Analysis)*. 26 trials were removed from the Parked task with a cutoff of 11.95s, and 25 trials from the Driving task with a cutoff of 12.9s.

As with Experiment 1, preliminary repeated measures ANOVAs were conducted to check whether Block interacted with the other factors in the Driving tasks (i.e. whether there were effects of practice/fatigue), for either the Distractor task (2 (Visibility) x 2 (Block)) and the Object-Detection task (3 (Viewport) x 2 (Visibility) x 2 (Block)). In all cases (and on all measures) there was a significant effect of Block, but importantly there were no interactions. Therefore it was deemed appropriate to collapse across Block to allow comparison across Parked vs. Driving tasks.

Results

Distractor Task Performance

A repeated measures 2 (Visibility) x 2 (Driving) ANOVA was conducted on $Comp_D$ scores (a log-transform of $Comp_D$ was applied to satisfy assumptions of normality since $Comp_D$ was positively skewed) to assess whether distractor task performance varied with level of difficulty of the other tasks. The ANOVA did not reveal a significant effect of Visibility $F\,(1, 25) = 2.67$, p=.11, $\eta^2 = .009$), but revealed an effect of Driving ($F\,(1, 25) = 67.66$, p<.001, $\eta^2 = .39$) and a significant interaction ($F\,(1, 25) = 6.35$, p=.018, $\eta^2 = .011$). The main effect of Driving suggests that individuals were less able to devote resources to the distraction task as the level of difficulty of other components of the task increased (steering and navigating as compared to being parked). Post-hoc Bonferroni comparisons showed that the interaction was due to higher (worse) $Comp_D$ scores for LV objects as compared to HV objects when Parked (p=.005) but not when Driving (p=1). One explanation for this interaction is that HV objects can be detected in the retinal periphery more readily than LV objects, which allows participants (whilst parked) to fixate (and foveate) the Distractor Task without significantly impairing their ability to detect HV objects (whereas LV objects need widespread gaze fixations to detect). The additional visual demands of driving removes this advantage because participants need to look in their direction of travel (Wilkie and Wann, 2003) rather than toward the Distractor task, and so the advantage for HV objects is lost.

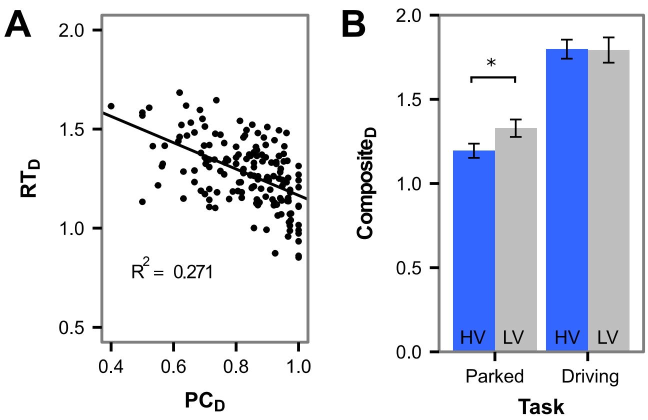

Figure 6 A) Percentage of odd-odd pairings correctly detected (PC_D) plotted against Reaction times for true positives (RT_D), per experiment run. B) $Comp_D$ when Parked (left) and Driving (right), for HV (Blue) and LV (Grey) objects. Lower $Comp_D$ scores indicates better performance on the distractor task. Error bars are standard error of the mean.

Object Detection Reaction Time Task Performance

RT_O for both Parked and Driving tasks are shown in Figure 7. A repeated measures 3 (Viewport) x 2 (Visibility) x 2 (Driving) ANOVA was assessed for systematic trends. The ANOVA revealed significant main effects of Viewport (F (2, 50) = 79.15, p<.001, η^2 = .26), Visibility (F (1, 25) = 282.61, p<.001, η^2 = .38), and Driving (F (1, 25) = 60.59, p<.001, η^2 = .15). There were also significant interactions between Viewport x Driving (F (2, 50) = 15.66, p<.001, η^2 = .06) and Visibility x Driving (F (1, 25) = 6.35, p=.018, η^2 = .018).

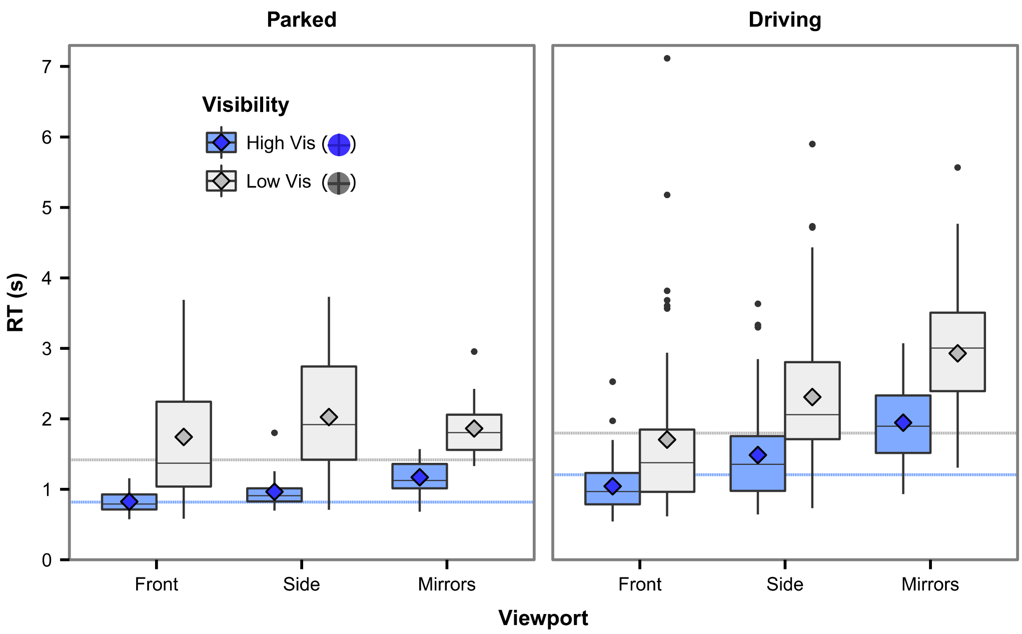

Figure 7 Boxplot showing distribution of RTs (seconds) across viewports (x-axis) and visibilities (Blue = high-visibility; Grey = low-visibility), for Parked (left panel) and Driving (right panel) tasks. Diamonds represent group means. Experiment 1 means (averaged across Viewports) are shown as dotted lines (Blue = high-visibility; Grey = low-visibility).

The Driving x Viewport interaction is plotted in Figure 8A. As per Experiment 1, driving did not cause a significant slowing of RTs for objects appearing through the front window ($p = .48$). However, Driving caused RTs to significantly slow for objects presented through the Side window ($p<.001$), and the largest slowing was for objects presented through the Mirrors ($p<.001$). The Driving x Visibility interaction is plotted in Figure 8B. It is clear that on *average* driving slowed RTs to a similar extent for both HV ($p<.001$) and LV objects ($p=.004$), shown by the diamonds on Fig8B. However, as with Experiment 1, there was considerably more individual variability in the extent which driving slowed RT_O to LV objects.

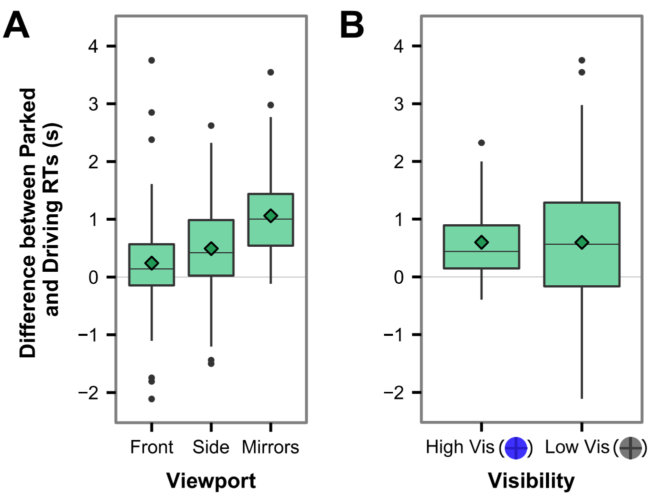

Figure 8 Investigating the Experiment 2 interactions between Driving and Viewport (A) and Driving and Visibility (B). Boxplots show the distributions of difference between means of Parked and Driving RTs. Diamonds represent mean differences.

Assessing the Effect of Cognitive Load

To assess whether the additional distraction task in Experiment 2 systematically altered patterns of responses (compared to Experiment 1), 'Experiment' was added as a between-subjects factor to a further ANOVA. Since there were no significant differences between HGV and Car drivers in Experiment 1 drivers were treated in this analysis as a single group. This allows a mixed model ANOVA with Experiment (E1: no Distractor vs. E2: added Distractor) as a between-subjects factor and Viewport, Visibility, and Driving as within-subjects factors. Significant main effects of Experiment (F (1, 65) = 25.76, p<.001, η^2 = .11), Visibility (F (1, 65) = 511, p<.001, η^2 = .34), Viewport (F (2, 130) = 183.16, p<.001, η^2 = .25), and Driving (F (1, 65) = 136.73, p<.001, η^2 = .15) were found, along with interactions between Visibility x Driving (F (1, 65) = 15.41, p<.001, η^2 = .017) and Viewport x Driving (F (1.58, 102.5) = 36.95, p<.001, η^2 = .061, ε =.79). Crucially, Experiment did not interact with any other factors, suggesting that although the Distractor task slowed RTs, this effect was consistent across the other variables.

Discussion

Experiment 2 examined whether the presence of an additional concurrent task (which required frequent glances to a VDU) interacted with patterns of RTs across viewports. It is clear that the additional cognitive load slowed RTs - the grand mean RT for Experiment 2 was 1.76s (SE=.047s), whereas the average RT for Experiment 1 was 1.38s (SE=.028s) – but apart from a general slowing the pattern of results across conditions were remarkably similar. Participants responded quickest to HV objects appearing in the front window when the driver was parked. The driving task exacerbated differences between viewports, and slowed responses to objects appearing on the side or the mirrors, but did not markedly slow responses to objects appearing in the front window.

It is perhaps surprising that the Distraction Task did not interact with the other task components, despite the visual (looking at the numbers) and cognitive (deciding on the appropriate response) loads entailed (Engström et al., 2005). One possible reason for the lack of interaction in the measured behaviours is that the compensatory behaviours were exhibited in ways that are currently unmeasured. Because the drivers were performing a number of concurrent tasks (detecting objects, responding to a VDU, steering, braking, and navigating) the participant could have ‘freed up’ resources for the distraction task by, for example, reducing lane variability or making fewer steering corrections (Engström et al., 2005; Kountouriotis et al., 2015). Because there were no explicit task instructions about driving performance (e.g. keep to the middle of the lane) we are unable to easily determine from the existing task whether such trade-offs took place.

An observed consequence of increased cognitive load is a greater concentration of gaze fixations on the road ahead (Recarte and Nunes, 2003; Victor et al., 2005), so it might have been predicted that participants would sample proportionally less from the (peripheral) mirrors during Experiment 2 than Experiment 1. The lack of Experiment x Viewport interaction suggests that this is not the case, and that the distribution of fixations across viewports were similar in both Experiments.

However, grouping the viewports as ‘mirrors’ ignores the differing visuo-spatial characteristics of each viewport, and it is not completely clear whether participants responded in a similar manner to individual mirrors or whether some mirrors were more rapidly responded to than others. Teasing out systematic differences between mirrors is not trivial since the ‘mirrors’ group consists of a selection of viewports differing in size and location relative to the driver. The next section examines this issue by determining whether RTs to objects in mirrors can be predicted simply by the differences in spatial position.

Further Analysis – Modelling Differences between Mirrors

The previous sections have established that drivers are generally quicker at responding to objects presented in the front window than to objects viewed in the mirrors. It remains unclear, however, what mechanism underpins the slowing of RTs to objects in the mirrors. This section considers whether there are properties that determine the systematic patterns of response times across mirrors.

We have demonstrated that RTs were quicker when objects were presented in front of the driver, presumably because people tend to look where they are going (Wilkie et al., 2010). Since larger eye-gaze movements take longer to execute, it might be predicted that the more eccentric a mirror is, the slower the RT would be.

The distribution of RTs to mirrors compared to the front window was captured by subtracting the front window RT from each mirror RT (per Visibility condition, per individual; RT_{Diff}). RT_{Diff} was taken (instead of the raw mirror RT) to account for varying intercepts across different individuals and object visibilities. This analysis was conducted solely on the RTs for the Driving task since there were not large differences between viewports in the Parked task. Since similar trends were observed in both experiments the data was pooled across them. To facilitate model comparison, both R^2 and Akaike information Criteria (Akaike, 1973) are reported. The R^2 value estimates the amount of outcome variance accounted for by the predictor so it is a useful and accessible way of assessing how well a given model fits the data. Because R^2 tends to increase with added parameters it can fall prey to overfitting (McElreath, 2016). AIC, however, gives an estimate of out-of-sample deviance so it avoids this problem, and is useful for comparing across models with different parameters.

To examine the extent that the distribution of RTs across mirrors could be explained by eccentricity from straight ahead we measured the distance (in radians) of the middle of each viewport from the point on the front windscreen which was in-line with the viewing position (Table 1). The first model, M1, assesses how well a single predictor – total eccentricity ($Total_{Ecc}$) – predicts RT_{Diff}. The fitted values are shown against RT_{Diff} in Figure 9 (crosses). The fit is remarkably good for a single parameter model, with $Total_{Ecc}$ explaining 48.4% of the variance in the slowing of RT responses (see Table 2). However, Figure 9 shows that $Total_{Ecc}$ predicted some viewports better than others. M1 assumes that as $Total_{Ecc}$ increases, RT_{Diff} will increase, but there are three viewports that do not fit this pattern. Class IV mirrors are more eccentric than Class II mirrors (see Table 1), so M1 predicts RT_{Diff} would be larger. Instead the opposite is true – people respond slightly quicker to Class IV than to Class II mirrors – so M1 overestimates RT_{Diff} for both Class IV mirrors. On the other hand, the Class VI mirror has the lowest $Total_{Ecc}$ but one of the highest RT_{Diffs}. M1 is unable to capture this behaviour and underestimates RT_{Diff} for Class VI.

The case of Class VI (short $Total_{Ecc}$ but high up in the visual field, and high RT_{Diff}) motivated splitting $Total_{Ecc}$ into two parameters – horizontal (Hz_{Ecc}) and vertical eccentricity (Vt_{Ecc}; Table 1). The added flexibility causes this second model (M2; Figure 9 filled triangles) to reduce, but not eliminate, the underestimation of RT_{Diff} for the Class VI mirror. However, M2 overestimates the RT_{Diff} for Class IV mirrors to a greater extent than M1, and does not improve estimates for any of the remaining mirrors. Indeed, Table 2 shows that despite the additional parameter M2 explains the same amount of variance as M1 and has a higher AIC (out-of-sample deviance) – so M1 should be preferred.

M2 was unable to correct for *both* the overestimation of Class IV mirrors and underestimation of the Class VI mirror because M2 predicts that Vt_{Ecc} would increase RTs, regardless of whether the direction of Vt_{Ecc} was low (Class IV) or high (Class VI) in the visual field. In order to resolve this conflict we added a sign to Vt_{Ecc} (Vt_{Sign}), so that Vt_{Sign} below the horizontal would be negative (so would *decrease* RT_{Diff}), and Vt_{Sign} above the horizontal would be positive (so would *increase* RT_{Diff}).

This third model – M3 – performs better than M1, explaining 51% of the variance and markedly reducing AIC (Table 3). Figure 9 (filled diamonds) shows that Vt_{Sign} gives M3 the additional flexibility needed to somewhat reduce underestimation of the Class VI mirror and reduce overestimation of Class IV mirrors. Whilst M3 still does not precisely estimate the exact RT_{Diff} for some individual viewports (in particularly Class VI is still underestimated), it does capture the broad pattern across mirrors fairly well. Appendix A examines whether the relationship between spatial position and response time varies with other factors, such as the presence of a distractor task.

Figure 9B shows how the fitted weights for Hz_{Ecc} and Vt_{Sign} relate to the extent of RT slowing with eccentricity. It is critical to note that the reported weights are specific to participants responding to an object of a constant optical size (of 1.95 degrees^2 as used in these experiments). These weights are likely to interact with object size, for example eccentricity might be weighted low when the object size is large (and therefore quickly discernible in the retinal periphery without large eye/head movements) but weighted high when the object size is small.

Table 1 Distance, in radians, of the middle of the mirror viewports measured from straight-ahead.

Mirror	*$Total_{Ecc}$ (Rads)*	*Hz_{Ecc} (Rads)*	*Vt_{Ecc} (Rads)*
IIR	.519	.519	.0
IIL	.732	.732	.0
IVR	.593	.519	(-).285
IVL	.786	.732	(-).285
V	.812	.785	.21
VI	.464	.326	.33

Table 2 Model comparison

Model	*F*	*df*	*P*	R^2	*AIC*
M1	767.7	1, 815	<.001	.484	2223.59
M2	381.5	2, 814	<.001	.483	2227.5
M3	424	2, 814	<.001	.51	2184.67

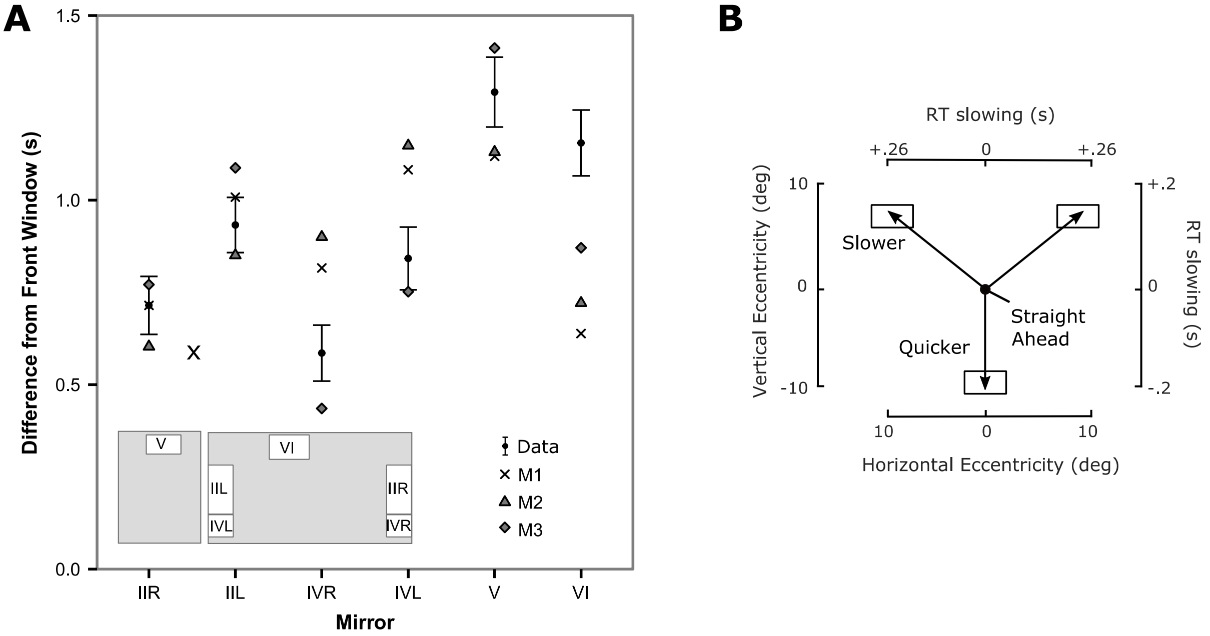

Figure 9. A) Modelling estimates shown against group RTs (with standard error bars), shown against fitted modelling values for M1 (crosses), M2 (filled triangles) and M3 (filled diamonds). Overlaid is a guide to mirror labelling. B) Schematic depicting the relationship between viewport eccentricity and reaction time slowing as predicted by model M3.

General Discussion

We have established that drivers were quickest when responding to objects that appeared in front of them, and slowest when objects appeared in mirrors that were far from the centre of the display (more eccentric). Closer inspection of the pattern of responses across mirrors revealed a systematic relationship between the extent of slowing and the eccentricity of the viewport: as horizontal eccentricity and vertical height of the viewport increased, so RTs to objects appearing there slowed.

These findings have important implications for road safety and vehicle design. Current cab designs rely on multiple mirrors in order to view regions around the cab to compensate for spatial blind-spots around the vehicle. Our results suggest that relying on a driver's ability to monitor the space around the vehicle indirectly (using mirrors) rather than directly (through the front window) is likely to lead to slower responses. There are some regions of space that cannot easily be seen directly (i.e. to the rear of the vehicle) and mirrors (or a suitable alternative such as a camera-view) will always likely to be needed to provide vision of these regions. However there are other blind-spot zones that lie in-front or to the side of the cab where there is the possibility of altering cab design to improve direct vision of these zones (Summerskill and Marshall, 2014). These regions are usually covered by the Class V & VI mirrors, but are problematic because the mirrors are high-up in the visual field and so produce slower responses (Figure 9B). Indeed slower responses to Class VI mirrors would help explain reports of front blind-spot accidents despite mirrors being set up correctly – perhaps the drivers simply did not see the pedestrians in time to respond appropriately (Cheng et al., 2016).

Our experiments set out to examine the fundamental perceptual-motor limitations of distributing visual search over multiple mirrors and windows. However, when generalising our findings to an HGV driver avoiding collisions in a real-world cab there are a number of factors which may moderate the relationship between eccentricity and reaction times proposed in Figure 9B. One potential issue is that whilst the simulation was purposely not matched to a single cab setup (bar meeting the minimum view requirements for mirrors), the viewport eccentricities used in the simulated setup meant that mirrors were closer to the centre of the display than would be the case in a real cab (Summerskill et al., 2015). Mirror positions with greater eccentricities would require even larger (so more time-consuming) eye and head movements, so we would predict even longer delays from drivers using real mirrors. In our experiments there was an equal probability of objects appearing in each viewport because the visual search task was deliberately designed to tackle the perceptual-motor challenges associated with keeping track of all the viewports. This required the driver to balance the visual demands on driving with performing well at the visual search task (which required fixations to be fairly equally distributed across viewports). Our findings suggest that there are limits to the extent individuals can simultaneously monitor multiple viewports. In the real-world, drivers are likely to try and compensate for these limits through learned scanning strategies which will direct visual search to particular viewports depending on the driving context (e.g. the Class VI mirror may be fixated frequently when a driver is pulling off from stationary, but looked at rarely when travelling at higher speeds). The extent that learned strategies interact with the relationship proposed in Figure 9B is a valuable avenue to explore for future research, and will need an experimental approach that attempts to reflect dynamic shifting of attention depending on context rather than an equal probability of a target object appearing in any viewport.

In the presented experiments one might also be concerned about the controlled nature and timing of target objects appearing across viewports: objects appeared suddenly without warning, potentially in any viewport (unlike seeing a pedestrian walking along the pavement and into the road where the driver might be able to predict the future real-world location). It should be highlighted that it is not uncommon for objects to be absent from a mirror one moment and then to be present the next (especially when monitoring six different mirrors), and so these were the visual conditions we were concerned with simulating. Whilst drivers were unable to predict where objects appeared this merely meant that gaze needed to be distributed fairly equally across all viewports to detect the objects. In the real-world it is likely that drivers would monitor mirrors less often than in our experiments but this would make driver reactions to suddenly appearing objects in mirrors even slower.

An additional way in which real-world mirrors may further slow driver RTs is the optical scaling that usually occurs (compressed images resulting due to the mirror increasing the field of view). In the main experiments reported here we ensured scaling was not a confounding factor by making all target objects of identical visual extent. To determine the impact of visual distortion on reactions we ran some additional driving conditions using a 'realistic' object that matched the visual characteristics of a pedestrian positioned in the world and so varied in optical size according to the scaling characteristics of the viewport in which it appeared (detail in Appendix B). Since objects appeared larger when viewed directly through the windows rather than through the mirrors, these conditions further increased the advantage for detecting an object directly through a window compared to indirectly via mirrors. Our results suggest that the time taken for the driver to execute an appropriate response would be significantly longer if a real object was viewed through a mirror placed eccentrically in the driver's field of view. Whilst it seems likely that the chances of collision with the object will be directly related to the ability of the driver to rapidly detect the object, further work is needed to establish the relationship between the RT differences observed here and actual changes in probability of collisions.

Related to the issue of learned driver responses, it is worth considering the nature of the response required from the drivers in these experiments. We deliberately chose a simple response action (pressing a paddle) that was independent of the driving controls to minimise potential confounds. However, in the real-world, the driver has to choose an appropriate response from large range of potential actions, which is likely to lead to slower reaction times than those observed in our experiments. Furthermore, the action selection process may interact with viewports in ways that are independent of eccentricity. For example, in the event of detecting a pedestrian in the Class VI blind-spot mirror the driver may be trained to brake suddenly without considering other options (due to the criticality of the situation and the limited number of options), whereas after detecting a cyclist in the Class V the driver could produce a more complex set of coordinated steering and braking responses that are selected from a wider range of actions (which may well lead to slower responses).

There are a number of potential implications of these findings depending on the driving environment being considered. Firstly, the faster responses to objects seen directly through the front window (compared to Class VI mirrors) will be of particular importance for HGVs when driving at slow speeds in urban environments (rather than at high speeds on motorways). Pedestrians walking in front of a stationary HGV are likely to be located in the front blind-spot. Even at slow driving speeds (5mph) the reduced RTs detecting a pedestrian in Class VI mirrors would lead to an increased stopping distance of ~1.5m (Arup and University of Leeds, 2017). There are a number of fatal incidences every year of involving these sorts of collisions, though it is often difficult to determine whether the driver did not use their mirrors, or whether detection was simply too slow. Irrespective of the cause, Transport for London are leading the way in implementing changes to government policy and recently (January 24th 2017) opened a consultation over establishing Direct Vision standards to quantify how appropriate an HGV cab design is for urban driving (Transport for London, 2017). By banning HGVs that are inappropriate for urban driving they aim to help to improve the chances of a drivers detecting vulnerable road users and thereby reducing road collisions.

The second main implication is for the future design of mirrors themselves. The location of mirrors have been driven historically by optical necessity – i.e. they are positioned so that the driver can see into a blind-spot from their seat position. This means that mirrors are often far from where the driver is looking (i.e. away from the road ahead) and require large, slow, head turns to check before making lane changes etc. Our findings indicate that detection would be quicker if the same (or better quality) visual information was provide lower in the scene and closer to the driver midline.

Many cabs are now surrounded by cameras that can capture similar images to mirrors but can be displayed in a viewport located anywhere in-vehicle. Our findings suggest that if the VDU was closer was closer to the driver's main line of sight then detection would be faster. Of course, it is impractical and unsafe to place a VDU on the front windscreen where it may obscure the road ahead. It might be tempting to place the VDU in a high-up position where it is less obtrusive, but we predict from the behaviours observed across both experiments that drivers would respond quicker to a VDU placed lower down in the scene (Figure 9B). Clearly these predictions need tested through controlled experiments. One particular issue with our modelling is that it assumes a linear relationship between eccentricity and reaction times, so Figure 9B depicts that the more the VDU is lowered, the quicker the driver responses. There are, however, other factors which may limit or eliminate the benefits of lowering a VDU. Firstly, the relationship between eccentricity and reaction times is likely to be non-linear: there may be little cost of eccentricity when drivers are able to make a rapid saccade to the VDU, but large costs when eccentricity is wide or low enough to require slower co-ordinated eye-head movements. Secondly, in the real-world the VDU is unlikely to be placed at the same depth plane as the front screen (as was the case in the simulation) so monitoring the VDU would require (slow) vergence changes. The best location for such a display therefore still needs to be empirically demonstrated.

An additional potential advantage of using cameras over mirrors is that the optical size of the image could be enhanced, therefore potentially leading to quicker detection rates (see Appendix B). There are a number of other issues with mirrors: they are often improperly adjusted, are unclean, or distort the image (Cook et al., 2011; Delmonte et al., 2012). Given these limitations it seems that we should certainly consider improving future cab design to replace mirrors with larger windows (to provide direct vision) and add smart VDUs that are positioned optimally to show the driver blind-spots that are relevant to their current driving manoeuvres.

Acknowledgements

The research was conducted under a contract with Arup as part of a larger project commissioned by Transport for London. TfL and Arup had a role in the initial scoping of the project, but no involvement in analysis, interpretation, or the manuscript write-up. An earlier version of the analyses is included in the end of project report (Arup, University of Leeds, 2017), but this data has been re-analysed and subject to further analyses in preparation for the current manuscript. We would like to thank Transport for London for funding the research, and colleagues at TfL (Hannah White, Glen Davies, & Vicky Sims) and Arup (Daniel Evanson, Rebecca Milner, & Harriet Western-Williams) for their insight during the project scoping exercise.

References

Ackery, A.D., McLellan, B., Redelmeier, D., 2012. Bicyclist deaths and striking vehicles in the USA. Inj. Prev. 18, 22–6. doi:10.1136/injuryprev-2011-040066

Akaike, H., 1973. Information theory and an extension of the maximum likelihood principle, in: Petrov, B.N., Csaki, F. (Eds.), Second International Symposium on Information Theory. pp. 267–281.

Arup, University of Leeds, 2017. Direct Vision vs Indirect Vision: A study exploring the potential improvements to road safety through expanding the HGV cab field of vision.

Bakeman, R., 2005. Recommended effect size statistics for repeated measures designs. Behav. Res. methods, 37, 379–384.

Bates, D., Maechler, M., Bolker, B., Walker, S., 2015. Fitting Linear Mixed-Effects Models using lme4. J. Stat. Softw. 67, 1–48. doi:doi:10.18637/jss.v067.i01.

Bíl, M., Bílová, M., Dobiáš, M., Andrášik, R., 2016. Circumstances and Causes of Fatal Cycling Crashes in the Czech Republic. Traffic Inj. Prev. 17, 394–399. doi:10.1080/15389588.2015.1094183

Cheng, Y.K., Wong, K.H., Tao, C.H., Tam, C.N., Tam, Y.Y., Tsang, C.N., 2016. Front blind spot crashes in Hong Kong. Forensic Sci. Int. 266, 102–108. doi:10.1016/j.forsciint.2016.05.013

Cook, S., Summerskill, S., Marshall, R., Richardson, J., Lawton, C., Grant, R., Bayer, S., Lenard, J., Clemo, K., 2011. The development of improvements to drivers' direct and indirect vision from vehicles - phase 2. Report for Department of Transport DfT TTS Project Ref: S0906/V8. Loughborough.

Davis, G., White, H., 2015. Reducing accidents between construction vehicles and cyclists. Civ. Eng. 168, 131–137. doi:10.1680/cien.14.00049

Deary, I.J., Liewald, D., Nissan, J., 2011. A free, easy-to-use, computer-based simple and four-choice reaction time programme: the Deary-Liewald reaction time task. Behav. Res. Methods 43, 258–68. doi:10.3758/s13428-010-0024-1

Delmonte, E., Manning, J., Helman, S., Basacik, D., Scoons, J., Chappell, J., Stannard, J., Jones, M., Knight, I., 2012. Construction logistics and cyclist safety. Tech. Rep.

Engström, J., Johansson, E., Östlund, J., 2005. Effects of visual and cognitive load in real and simulated motorway driving. Transp. Res. Part F Traffic Psychol. Behav. 8, 97–120. doi:10.1016/j.trf.2005.04.012

Engström, J., Markkula, G., Victor, T., Merat, N., 2017. Effects of cognitive load on driving performance : The cognitive control hypothesis, Human factors. *in press*.

Gilbert, K., McCarthy, M., 1994. Deaths of cyclists in London 1985-92: the hazards of road traffic. BMJ. 308, 1534. doi:https://doi.org/10.1136/bmj.308.6943.1534

Goldring, J.E., Dorris, M.C., Corneil, B.D., Ballantyne, P. a, Munoz, D.P., 1996. Combined eye-head gaze shifts to visual and auditory targets in humans. Exp. Brain Res. 111, 68–78. doi:10.1007/BF00229557

Green, M., 2000. "How Long Does It Take to Stop?" Methodological Analysis of Driver Perception-Brake Times. Transp. Hum. Factors, 195–216. doi:10.1207/STHF0203_1

Hines, T.M., 1990. An odd effect: Lengthened reaction times for judgments about odd digits. Mem. Cognit. 18, 40–46. doi:10.3758/BF03202644

Horrey, W., Wickens, C., 2007. In-Vehicle Glance Duration: Distributions, Tails, and Model of Crash Risk. Transp. Res. Rec. J. Transp. Res. Board. 22–28. doi:10.3141/2018-04

Jansen, R.J., Van Egmond, R., De Ridder, H., 2016. Task prioritization in dual-tasking: Instructions versus preferences. PLoS One. 11, 1–27. doi:10.1371/journal.pone.0158511

Keigan, M., Cuerden, R., Wheeler, A., 2009. Analysis of police collision files for pedal cyclist fatalities in London, 2001-2006. TRL Report. R438.

Kountouriotis, G.K., Wilkie, R.M., Gardner, P.H., Merat, N., 2015. Looking and thinking when driving: The impact of gaze and cognitive load on steering. Transp. Res. Part F Traffic Psychol. Behav. 34, 108–121. doi:10.1016/j.trf.2015.07.012

Lawrence, M.A., 2016. ez: Easy Analysis and Visualization of Factorial Experiments. R package version 4.4-0.

Leigh, J., Zee, D.S., 2006. The Neurology of Eye Movements. Oxford University Press.

Levy, J., Pashler, H.E., Boer, E.R., 2006. Is there any stopping the psychological refractory period? Psychol. Sci. 17, 228–234. doi:10.1111/j.1467-9280.2006.01690.x

Luce, R., 1986. Response Times. Oxford University Press, New York.

Manson, J., Cooper, S., West, A., Foster, E., Cole, E., Tai, N.R.M., 2012. Major trauma and urban cyclists: physiological status and injury profile. Emerg. Med. J. 30, 32–37. doi:10.1136/emermed-2011-200966

McElreath, R., 2016. Statistical Rethinking: A Bayesian Course with examples in R and Stan. CRC Press, Boca Raton, FL.

Niewoehner, W., Berg, A., 2005. Endangerment of Pedestrians and Bicyclists At Intersections 1–15.

Osaka, N., 1976. Visual Reaction Time as a Function of Target Size and Retinal Eccentricity in the Peripheral Visual Field. Jpn. Psychol. Res. 18, 183–190.

Pattinson, W., Thompson, R.G., 2014. Trucks and Bikes: Sharing the Roads. Procedia - Soc. Behav. Sci. 125, 251–261. doi:10.1016/j.sbspro.2014.01.1471

R Core Team, 2016. R: A language and Environment for Statistical Computing.

Ratcliff, R., 1993. Methods for dealing with reaction time outliers. Psychol. Bull. 114, 510–532. doi:10.1037/0033-2909.114.3.510

Recarte, M.A., Nunes, L.M., 2003. Mental workload while driving: effects on visual search, discrimination, and decision making. J. Exp. Psychol. Appl. 9, 119–37. doi:10.1037/1076-898X.9.2.119

Sareen, P., Ehinger, K. a, Wolfe, J.M., 2015. Through the looking-glass: Objects in the mirror are less real. Psychon. Bull. Rev. 980–986. doi:10.3758/s13423-014-0761-8

Simonen, R. L., Battie, M. C., Videmen, T. & Gibbons, L. E., 1995. Comparison of Foot and Hand Reaction Times among Men: A Methologic Study Using Simple and Multiple-Choice Repeated Measurements. Perc. Motor. Skills. 80, 1243-1249.

Sodhi, M., Reimer, B., Llamazares, I., 2002. Glance analysis of driver eye movements to evaluate distraction. Behav. Res. Methods. Instrum. Comput. 34, 529–38. doi:10.3758/BF03195482

Summala, H., Lamble, D., Laakso, M., 1998. Driving experience and perception of the lead car's braking when looking at in-car targets. Accid. Anal. Prev. 30, 401–407. doi:10.1016/S0001-4575(98)00005-0

Summerskill, S., Marshall, R., 2014. The design of category N3 vehicles for improved driver direct vision. Technical Report. 0–118.

Summerskill, S., Marshall, R., Paterson, A., Reed, S., 2015. Understanding Direct and Indirect Vision in Heavy Goods Vehicles. Technical Report.

T&E, 2016. Eliminating truck blind spots – a matter of (direct) vision.

Traffic Analysis Centre, 2012. Traffic levels on major roads in Greater London 1993-2010.

Transport for London, 2017. Direct Vision Standard Phase 1 Consultation support document.

Transport for London, 2015. Travel in London, Report 8.

Victor, T.W., Harbluk, J.L., Engström, J.A., 2005. Sensitivity of eye-movement measures to in-vehicle task difficulty. Transp. Res. Part F Traffic Psychol. Behav. 8, 167–190. doi:10.1016/j.trf.2005.04.014

Whelan, R., 2008. Effective analysis of reaction time data. Psychol. Rec. 58, 475–482.

White, A.D., Mushtaq, F., Raw, R.K., Giles, O.T., Crook, I.C., Tomlinson, J., Miskovic, D., Lodge, J.P.A., Wilkie, R.M., Mon-Williams, M., 2016. Does monitor position influence visual-motor performance during minimally invasive surgery? J. Surg. Simul. 3, 1–7. doi:10.1102/2051-7726.2016.0001

Wierwille, W.W., 1993. Visual and Manual Demands of In-Car Controls and Displays, in: Peacock, B., Karwowski, W. (Eds.), Automotive Ergonomics. Taylor & Francis, Washington DC, pp. 299–320.

Wilkie, R.M., Kountouriotis, G.K., Merat, N., Wann, J.P., 2010. Using vision to control locomotion: looking where you want to go. Exp. Brain Res. 204, 539–47. doi:10.1007/s00221-010-2321-4

Wilkie, R.M., Wann, J.P., 2003. Eye-movements aid the control of locomotion. J. Vis. 3, 677–684. doi:10.1167/3.11.3

Woolsgrove, C., 2014. ECF report on HGV cabs direct vision and amendments to Directive 96 / 53.

Yang, Q., Bucci, M.P., Kapuola, S., 2002. The latency of sccades, vergence, and combined eye movements in children and in adults. Investig. Ophthalmol. Vis. Sci. 43, 2939–2949.

Appendix

A. Further Modelling: Varying effects to examine interactions

Whilst model M3 provides the most parsimonious explanatory account for variance in performance across conditions there are other factors that can be considered, including: the added load of the Distraction Task (Experiment), whether the participant was a Car or HGV driver (Driver), or visibility of the object (Visibility). Interactions were tested by adjusting M3 to allow predictor weights to vary across other factors (random effects). The predicted weights are reported in Table A.1 alongside AIC, which is useful for comparing across models with different numbers of parameters (the AIC value for M3 was 2184.67). Crucially, having predictor weights that vary across Experiment or Driver did not improve the model's out-of-sample deviance, suggesting that the relationship between spatial position and response times is not dependent on concurrent tasks (Experiment), or prior experience with multiple mirrors (Driver). Allowing effects to varying across Visibility led to a slight reduction in AIC (i.e. better performance). This is evidence for an interaction with Visibility (AIC=2181.92) insofar that Hz_{Ecc} had a greater detrimental impact on RT_{Diff} when responding to LV rather than HV objects.

Table A.1 Fitted predictor weights and AIC values for random effects models tested.

Model	Level	Predictor *β*		*AIC*
		Hz_{Ecc}	Vt_{Sign}	
M3a (Driver)	Car	1.56	1.24	.2192.21
	HGV	1.07	.6	
M3b (Experiment)	Exp1	1.33	.99	2193.74
	Exp2	1.71	1.37	
M3c (Visibility)	HV	1.22	1.16	2181.92
	LV	1.74	1.12	

B. Detecting Pedestrians of Varying Optical Size

To examine the effect of optical size on detection times we reran the Driving task with a pedestrian object placed in the real-world (rather than on the viewport screen) at a constant distance from the viewpoint (5m) so that the optical size of the pedestrian scaled commensurately with the characteristics of the viewport it was presented in. This means that the pedestrian was largest when viewed through the Front and Side windows, smaller in the Class II mirrors, smaller still in Class V and VI mirrors, and smallest in Class IV mirrors. All participants took part in this version of the task. The group RTs are shown in Figure A.1A. A 2 (Experiment) x 3 (Viewport) mixed model ANOVA was conducted to assess for systematic trends. Significant main effects of Experiment ($F\,(1, 65) = 20.64$, $p<.001$, $\eta^2 = .016$) and Viewport were ($F\,(1.63, 106) = 358.17$, $p<.001$, $\eta^2 = .68$, $\varepsilon = .82$), but there was not a significant interaction.

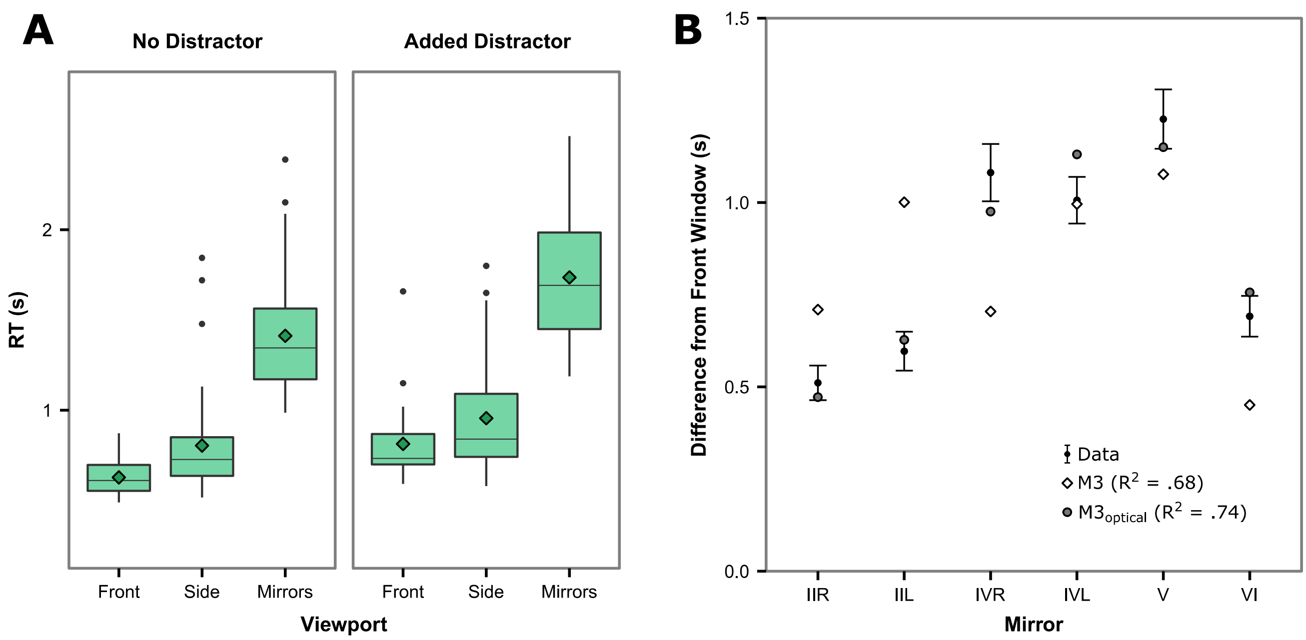

Figure A.1. A) Group RTs for Pedestrian objects for Experiment 1 (No Distractor) and Experiment 2 (Added Distractor). B) Modelling estimates for M3 and $M3_{optical}$ shown alongside group RT_{Diff}, with bars representing standard error of the mean.

We were particularly interested in whether the extent of slowing of responses in mirrors could be predicted by viewport eccentricities, as was the case with objects of a constant optical size, or whether detection rates were also determined by the optical size of the target object. Figure A.1B shows that M3 ($RT_{Diff} = Hz_{Ecc} + Vt_{Sign}$) does fairly well at capturing the patterns of responses across mirrors, explaining 68% of the variance in mirror RT_{Diff} (AIC = 730.14), although there are some mirrors (IIL & IVR) M3 does not capture well. Adding an optical scaling parameter (i.e. how many times smaller the object appears in a viewport compared to the front window) increases the variance explained to 74% and also greatly reduces out-of-sample deviance (AIC = 648.08). It is clear from Figure A.1B that $M3_{optical}$ ($RT_{Diff} = Hz_{Ecc} + Vt_{Sign} +$ Optical-Scaling) does remarkably well at capturing each individual mirror, suggesting that the optical size of an object influences the relationship between viewport eccentricity and reaction times.