Short Research Article

Eye Movements Reveal How Readers Infer Intentions From the Beliefs and Desires of Others

Matthew Haigh[1] and Jean-François Bonnefon[2]

[1]Department of Psychology, Northumbria University, UK,
[2]Centre National de la Recherche Scientifique, Toulouse, France

Abstract. We examine how the beliefs and desires of a protagonist are used by readers to predict their intentions as a narrative vignette unfolds. Eye movement measures revealed that readers rapidly inferred an intention when the protagonist desired an outcome, even when this inference was not licensed by the protagonist's belief state. Reading was immediately disrupted when participants encountered a described action that contradicted this inference. During intermediate processing, desire inferences were moderated by the protagonist's belief state. Effects that emerged later in the text were again driven solely by the protagonist's desires. These data suggest that desire-based inferences are initially drawn irrespective of belief state, but are then quickly inhibited if not licensed by relevant beliefs. This inhibition of desire-based inferences may be an effortful process as it was not systematically sustained in later steps of processing.

Keywords: theory of mind, narrative comprehension, conditional reasoning, eye tracking

People can infer the intentions of other individuals by tracking their beliefs and desires. This *theory of mind* (ToM) plays an important role in helping people navigate everyday social situations. A central issue for ToM research is precisely how fast and automatic this reasoning can be. One basic consideration is that ToM needs to keep pace with the situations it applies to. If people rely on ToM to make inferences in fast moving social interactions, ToM must be fast, too (Apperly, 2012). Consistent with this assumption, studies using process tracking techniques with high temporal resolution have shown that adults rapidly and routinely make inferences based on the beliefs (e.g., van der Wel, Sebanz, & Knoblich, 2014) and desires (e.g., Ferguson & Breheny, 2011) of others.

While previous research has tended to examine these belief and desire-based inferences in isolation, everyday social or narrative situations can yield simultaneous inferences with often conflicting conclusions (see Malle & Holbrook, 2012). For example, a protagonist may desire an outcome, but at the same time not hold the relevant knowledge to achieve this outcome (note that we use the terms "belief" and "knowledge" interchangeably). In this article, we use eye movement measures of narrative comprehension to examine how these belief and desire states are used by readers to predict the intentions of a fictional protagonist as a story unfolds.

Evidence that people draw inferences about the likely intentions and behaviors of others comes from a variety of fields including developmental, social and cognitive psychology. These inferences can be informed by the perceived mental state of an individual (e.g., their specific beliefs and desires) or by more universal folk understandings of human motivation (e.g., Malle, 1999; Miller, 1999). For example, we know from reasoning studies that adults consistently expect protagonists to do whatever is in their best interest (Bonnefon, 2009). When given the conditional sentence (1), without any additional context, reasoners consistently infer that Alice is going to renew her insurance over the Internet (e.g., Bonnefon, 2012; Bonnefon, Haigh, & Stewart, 2013; Bonnefon & Hilton, 2004; Egan & Byrne, 2012; Evans, Neilens, Handley, & Over, 2008):

(1) If Alice renews her insurance over the Internet, she will save £100

Such inferences draw on assumptions about the most likely beliefs and desires of the protagonist. If told that Alice did not in fact renew her insurance online, reasoners conclude that she did not know about the Internet deal, or that she did not care about saving £100 (Bonnefon, Girotto, & Legrenzi, 2012; see also Malle & Knobe, 1997). In sum, when given a conditional rule like (1), reasoners assume that the protagonist knows about the rule (belief), cares about its outcome (desire), and they then derive a ToM inference about what the character will do (intention).

These inferred beliefs, desires, and intentions play a central role in how we simulate the mental states of others, and they are closely monitored as other people are either viewed (e.g., Schneider, Bayliss, Becker, & Dux, 2012) or

DOI: 10.1027/1618-3169/a000290

described in a narrative text (e.g., Zwaan, Langston, & Graesser, 1995).

Narrative comprehension studies have shown that readers routinely monitor the desires of fictional characters and expect them to act in a manner that fulfills these desires. In a line-by-line self-paced reading experiment, Huitema, Dopkins, Klin, and Myers (1993) presented participants with a context in which the protagonist's desire was explicitly stated "...[Dick] knew he wanted to go somewhere warm where he could swim and sunbathe." A later sentence then described an action that was either consistent or inconsistent with this desire ("He phoned TWA and booked a ticket to Florida/Alaska"). Consistent with the hypothesis that readers track desires and expect protagonist's to act accordingly, readers took longer to process this line when the action mismatched the desire. This effect has been replicated even when the protagonist's desire is merely implied, rather than explicitly stated (e.g., "[Dick] had always been a real sun-worshipper"; Poynor & Morris, 2003). More recent research using the visual world paradigm has shown that the higher order desires of a protagonist cause listeners to make rapid predictive inferences as a narrative situation unfolds in real time, even when these desires conflict with the protagonist's basic preferences (Ferguson & Breheny, 2011).

In addition, there is converging evidence that adults can track the beliefs of other agents without the need for effortful, strategic processing. Unlike young children, adults tend to correctly attribute true and false beliefs to others and use these attributions to predict their actions. More to the point, there is strong evidence that adults can make these belief-based inferences unconsciously and automatically (Kovács, Téglás, & Endress, 2010; Qureshi, Apperly, & Samson, 2010; Samson, Apperly, Braithwaite, Andrews, & Bodley Scott, 2010; Schneider et al., 2012; van der et al., 2014) including during narrative comprehension (de Vega, Díaz, & León, 1997). However, adult ToM reasoning is not always faultless, with evidence from both offline (e.g., Birch & Bloom, 2007; Keysar, 1994) and online tasks (e.g., Ferguson & Breheny, 2012; Keysar, Barr, Balin, & Brauner, 2000) showing that the attribution of beliefs to others can be subject to interference from an observer's privileged knowledge. Specifically, observers are susceptible to erroneously attributing their own beliefs (knowledge) to actors that do not possess these beliefs.

Although the evidence described above suggests that both desire- and belief-based inferences are drawn routinely and rapidly by adults, they do not reveal how the perceived beliefs and desires of an actor interact in the mind of an observer to predict an intention. When belief and desire states simultaneously or serially produce conflicting intent inferences, an observer must rapidly prioritize just one. Below we consider two accounts of how this process may unfold during narrative comprehension. The first is based on the necessity of both beliefs and desires to infer intent (Malle & Knobe, 1997) and the second is based on their relative salience in the mind of an observer (Malle & Holbrook, 2012).

First, there are solid normative and empirical grounds to predict that an agent's belief (or desire) state will always moderate inferences based on their desires (or beliefs), due to the fundamental assumption that neither beliefs nor desires are sufficient in their own right to reliably infer an intention. Malle and Knobe (1997) found that the majority of participants (81%) inferred intent when a protagonist desired an outcome and held the relevant beliefs to make this happen, but not when information about one of these factors was absent (only 21% inferred intent when there was desire but no belief and 31% when there was belief but no desire). For example, if a protagonist does not hold the belief that an action will lead to a desired outcome (e.g., Alice *does not know* that if she renews her insurance over the Internet, she will save £100), then there are no grounds for an observer to make a desire-based inference about her likely behavior (e.g., that she will renew over Internet), no matter how much she desires the outcome (saving £100). Likewise, if a protagonist does not desire an outcome, there are no grounds for a belief-based inference. In other words, an intent inference based on desires (or beliefs) should *only* be made when it is licensed by the protagonist's belief (or desire) state.

However, recent research has shown that the beliefs and desires of others are not necessarily equivalent in the mind of an observer. Consistent with developmental evidence that desire-based inferences develop much earlier than belief-based inferences (Wellman & Woolley, 1990) it has been shown that adult observers perceive the intentions and desires of other people to be consistently more salient than beliefs (Malle & Holbrook, 2012). These data suggest that the desires of others are more likely to be at the forefront of an observer's mind as a social situation unfolds. If this is the case, then desire-based inferences may be drawn even in situations where they are not licensed by the protagonist's beliefs (e.g., inferring that *Alice* will renew over the Internet even though she does not know the benefits of this action). This would be consistent with a growing body of evidence that the privileged beliefs of an observer can interfere with their reasoning about the behavior of others; as demonstrated by the so-called "*curse of knowledge*" effect (e.g., Birch & Bloom, 2007) and the related tendency for adults to use "egocentric" heuristic strategies (e.g., Keysar et al., 2000).

Experiment

We tested our two accounts of how readers infer intentions in an eye tracking experiment where participants read vignettes describing the beliefs, desires, and actions of a protagonist. This method allows us to determine which account best explains the data at different temporal stages of the comprehension process. Consider the following example:

(2) a. Alice needed to renew her car insurance before it expired.
b. She knew that if she renewed over the Internet she would save £100. [Belief]
c. Such a saving was important as she was struggling financially and desperately needed to save money. [Desire]
d. After gathering together the relevant documents she renewed her policy over the phone. [Action]
e. The call lasted nearly half an hour.

Account 1: *Intent inferences occur when the protagonist holds relevant beliefs and desires* (Malle & Knobe, 1997).

Sentences (2b) and (2c) establish belief and a desire states that would typically lead reasoners to infer that Alice will renew her car insurance over the Internet (Bonnefon & Hilton, 2004). If readers keep track of the belief and desire states described by (2b) and (2c), then they should rapidly detect that sentence (2d) is anomalous (Huitema et al., 1993). Detection should be manifest on the verb phrase at the end of this sentence (i.e., "renewed her policy over the phone") in the form of immediate disruption to normal left-to-right eye movements, an increase in fixation times and more regressive saccades back into this region. Residual or delayed effects may also emerge on the following sentence (2e) (Rayner, 1998). If *both* beliefs and desires are required to predict intent, then the same sentence should not be perceived as anomalous when Alice does not desire the outcome (e.g., sentence (2c) replaced with "because she was very wealthy such a saving was not important to her").

Now consider the same vignette in which the belief described in sentence (2b) is replaced with "She didn't know that if she renewed over the Internet she would save £100." In this instance there is no reason to infer that Alice will favor any one method of renewal, although sentence (2c) continues to state her strong desire to save money. If beliefs are required to licence desires, then readers should not make a desire-based intent inference (i.e., they do not predict that Alice will renew over the Internet) because the protagonist's beliefs give no good reason to favor any one course of action. In this instance eye movements and fixation times to sentence (2d) should not differ between a condition in which the protagonist desires the outcome and a condition in which the protagonist does not desire the outcome.

Account 2: *Intent inferences occur when the protagonist holds relevant desires* (Malle & Holbrook, 2012).

An alternative possibility is that desire-based inferences are drawn whenever the protagonist holds a relevant desire (irrespective of their beliefs). If this is the case, then readers will always predict that Alice will renew her insurance over the Internet in situations where she desires saving money. This inference is drawn regardless of what the reader knows about the protagonist's beliefs, due to interference from their privileged knowledge (i.e., the reader knows that renewing over the Internet will achieve her desires) (e.g., Birch & Bloom, 2007). In this instance (2d) should cause disruption to eye movements in conditions where the protagonist desires the outcome, relative to conditions where the protagonist does not desire the outcome. Delayed or residual effects may also be observed on sentence (2e).

Method

Participants

Thirty-six native English speakers completed the experiment (26 female, mean age 24.7, normal vision, no diagnosed language impairment). Seven additional participants did not complete the experiment because they consistently struggled to fixate on the gaze trigger and their data were excluded. Participants were recruited from Northumbria University and compensated with £6 cash.

Design and Materials

Experimental items were vignettes describing the behavior of a protagonist. In each vignette the protagonist had the opportunity to take an action that would result in a personal benefit, but ultimately failed to take that action. Importantly, prior context manipulated whether or not the protagonist knew about the potential benefit of the action (belief), and whether or not the protagonist attached subjective importance to this benefit (desire). This resulted in a fully crossed 2 (believed vs. not believed) × 2 (desired vs. not desired) repeated measures design.

Twenty-four experimental scenarios were created, each with four versions. The 96 resulting vignettes were five sentences long (see Table 1 for a detailed example and Electronic Supplementary Material ESM 1 for a complete list). Sentence 1 provided contextual information and did not differ across conditions. Sentence 2 stated a conditional rule of the form *if action then benefit* that the protagonist either knew or did not know (belief). Sentence 3 was manipulated so that the benefit was either subjectively desired or undesired by the protagonist (desire). Sentence 4 described the protagonist failing to take the action identified in the conditional rule. Sentence 5 provided a continuation of the narrative. Within each scenario, sentences 4 and 5 were lexically identical across all conditions.

One version of each scenario was assigned to one of four Latin squared presentation lists. Nine participants were assigned to each list. The 24 target items were interspersed with 48 unrelated items and presented in a different random order to each participant. Comprehension questions requiring a "yes" or "no" response followed 50% of items, and were solved with a mean accuracy of 93.4%.

Procedure

Participants were simply instructed to read the vignettes silently for comprehension. Eye movements were recorded using an Eyelink 1,000 in desktop mount configuration. Viewing was binocular and recordings were sampled from the right eye at 1,000 Hz. Vignettes were presented in size 20 Arial font on a CRT monitor 80 cm from the participants' eyes. The head was stabilized using a chin rest.

The eye tracker was first calibrated using nine fixation points. Each trial began with a gaze trigger corresponding to the position of the first letter. Fixation on the gaze trigger caused the vignette to appear. After reading a vignette participants pressed a button on a handheld controller to advance. The session lasted approximately 30–40 min and all participants were given a short break at halfway.

Table 1. Example of the four experimental conditions. Manipulated context appears in italics, and the critical analysis region is highlighted in bold. The critical region was always the final verb phrase of sentence four, which is lexically identical across the four conditions

Believed	Desired	Alice needed to renew her car insurance before it expired. *She knew* that if she renewed over the Internet she would save £100. *Such a saving was important as she was struggling financially and desperately needed to save money.* After gathering together the relevant documents she |**renewed her policy over the phone.** CRITICAL| The call lasted nearly half an hour. POST–CRITICAL|
	Not desired	Alice needed to renew her car insurance before it expired. *She knew* that if she renewed over the Internet she would save £100. *However, because she was very wealthy such a saving was not important to her.* After gathering together the relevant documents she |**renewed her policy over the phone.** CRITICAL| The call lasted nearly half an hour. POST–CRITICAL|
Not believed	Desired	Alice needed to renew her car insurance before it expired. *She didn't know* that if she renewed over the Internet she would save £100. *Such a saving was important as she was struggling financially and desperately needed to save money.* After gathering together the relevant documents she |**renewed her policy over the phone.** CRITICAL| The call lasted nearly half an hour. POST–CRITICAL|
	Not desired	Alice needed to renew her car insurance before it expired. *She didn't know* that if she renewed over the Internet she would save £100. *However, because she was very wealthy such a saving was not important to her.* After gathering together the relevant documents she |**renewed her policy over the phone.** CRITICAL| The call lasted nearly half an hour. POST–CRITICAL|

Analysis

We analyzed the critical region (the verb phrase in sentence 4) and the post-critical region of text (sentence 5), see Table 1. These regions were lexically identical within items across the four conditions. Fixations < 80 ms were pooled with adjacent fixations, while fixations < 40 ms were excluded if they were not within three characters of another fixation. Fixations > 1,200 ms were truncated. The two analysis regions were examined using four processing measures (see below). We do not report *first pass* reading times as this measure produces unreliable results in situations where the region of interest includes more than one word *and* the percentage of First Pass Regressions Out differs between conditions (Pickering, Frisson, McElree, & Traxler, 2004).

First pass regressions out is the percentage of trials in which a first pass fixation in a region is followed by a fixation to an earlier region. This measure indicates the degree to which left to right eye movements are disrupted when first encountering a region of text.

Regression path reading time is the summed duration (in ms) of fixations to a region of text, from *first entering* to *first exiting to the right* (not including time spent outside the region to the left). It cannot be reliably computed for the post-critical region because there is no text to the right of this sentence.

Regressions in is the percentage of trials in which a fixation to a region was preceded by a fixation to a later region. This indicates the degree to which a region of text is revisited. It cannot be computed for the post-critical region as this is the final sentence.

Total time is the summed duration (in ms) of *all fixations* to a region (including advance peeking and revisits).

Analyses were conducted using 2 × 2 (Belief × Desire) repeated measures ANOVAs with subjects (F_1) and items (F_2) as random factors. Interactions were further analyzed using paired *t*-tests. Eye movement comparisons were always between lexically identical regions of text.

Results

Table 2 displays the mean and standard error of the mean of our five measures, averaged across subjects, for all conditions and regions of text (see Electronic Supplementary Material ESM2 for the individual subject and item means). Table 3 displays the results of the ANOVAs by subjects and items.

As the critical region was read participants were clearly sensitive to the protagonist's desires on measures of first pass regressions out, regression path and total time. Early regressive saccades out of the critical region (First Pass Regressions Out) are influenced only by the protagonist's desires (irrespective of their beliefs). However, significant interaction effects show that the time taken to then initially go past this region (Regression Path) and the Total Time taken to read this region *are* moderated by the beliefs of the protagonist. Regression path reading time to the critical region was impacted by desire *only* when the protagonist knew how to achieve her desire (888 ms vs. 759 ms; $t_1(35) = 3.58$, $p = .001$, $d = .66$; $t_2(23) = 3.04$, $p = .01$, $d = .76$), but not in the absence of this belief (876 ms vs. 845 ms; $t_1(35) = .72$, $p = .48$, $d = .12$; $t_2(23) = .79$, $p = .44$, $d = .17$). Similarly, total reading time to the critical region was strongly impacted by desire when the protagonist knew how to achieve her desire (1,048 ms vs. 823 ms; $t_1(35) = 5.41$, $p < .001$, $d = 1.12$; $t_2(23) = 5.21$, $p = < .001$, $d = 1.37$), but the effect size was much smaller in the absence of this belief (1,023 ms vs. 929 ms; $t_1(35) = 2.01$, $p = .05$, $d = .36$; $t_2(23) = 2.62$, $p = .02$, $d = .57$).

Table 2. Mean reading times and regressions by analysis region and condition (means averaged over subjects, standard error of the mean in parenthesis)

	First pass regressions out (%)	Regression path (ms)†	Regressions in (%)†	Total time (ms)
Critical region				
1. Believed/Desired	22.5 (3.0)	888 (46)	20.5 (3.1)	1,048 (58)
2. Believed/Not desired	13.9 (2.8)	759 (29)	15.9 (2.5)	823 (31)
3. Not believed/Desired	20.4 (3.3)	876 (46)	22.3 (3.1)	1,023 (54)
4. Not believed/Not desired	15.3 (3.0)	845 (39)	16.3 (2.4)	929 (43)
Post-critical region				
1. Believed/Desired	54.8 (4.7)	n/a	n/a	1,815 (72)
2. Believed/Not desired	39.9 (4.8)	n/a	n/a	1,704 (75)
3. Not believed/Desired	50.1 (4.1)	n/a	n/a	1,838 (80)
4. Not believed/Not desired	45.0 (4.6)	n/a	n/a	1,837 (75)

Note. †Regression Path and Regressions In are not applicable to the post-critical region as this was always the final sentence of the vignette.

Table 3. Analysis of variance results for each measure of interest (significant effects in bold)

			By participants			By items		
Region	Measure	Predictor	$F_1(1, 35)$	p	η_p^2	$F_2(1, 23)$	p	η_p^2
Critical	First pass regression out	Belief	< 1			< 1		
		Desire	**7.02**	**.01**	**.17**	**9.22**	**.01**	**.29**
		Interaction	< 1			< 1		
	Regression path time	Belief	1.74	.20	.05	1.33	.26	.06
		Desire	**6.73**	**.01**	**.16**	**7.57**	**.01**	**.25**
		Interaction	**3.88**	**.057**	**.10**	**5.14**	**.03**	**.18**
	Regressions in	Belief	< 1			< 1		
		Desire	**5.73**	**.02**	**.14**	**4.65**	**.04**	**.17**
		Interaction	< 1			< 1		
	Total time	Belief	< 1			1.69	.21	.07
		Desire	**26.38**	**< .001**	**.43**	**32.91**	**< .001**	**.59**
		Interaction	**4.70**	**.04**	**.12**	**5.51**	**.03**	**.19**
Post-critical	First pass regression out	Belief	< 1			< 1		
		Desire	**8.48**	**.01**	**.20**	**4.82**	**.04**	**.17**
		Interaction	2.65	.11	.07	2.68	.11	.10
	Total time	Belief	2.71	.11	.07	4.54	.04	.16
		Desire	1.10	.30	.03	2.82	.11	.11
		Interaction	2.67	.11	.07	2.06	.17	.08

On later measures of processing we found a main effect of desire, with first pass fixations to the post-critical region more likely to be followed by a revisit to an earlier region when the protagonist desired the outcome (52%) relative to when they did not desire the outcome (42%). Likewise, there were more return visits to the critical region when the protagonist desired the outcome than when they did not desire the outcome (21% vs. 16% regressions in to critical region). These effects were not moderated by the protagonist's beliefs.

Discussion

In the early stages of integrative processing there was clear evidence that readers made rapid intent inferences based on the desires of the protagonist. We set out to test two accounts that specified the necessary conditions for such inference generation to occur. The first account predicted that desire-based inferences would only be made when the protagonist explicitly knew how to achieve their desire (cf. Malle & Knobe, 1997). The second account predicted that desire-based inferences would be made regardless of the protagonist's belief state (cf. Malle & Holbrook, 2012). We found evidence that these accounts explain different temporal stages of the comprehension process.

On the earliest measure of processing (First Pass Regressions Out of the critical region) there was rapid disruption to normal fluent reading when the protagonist desired an outcome but acted in a manner that failed to realize this outcome. This early effect occurred irrespective of the protagonist's beliefs and is consistent with an account in

which the desires of the protagonist are more salient than beliefs in the mind of a reader (Malle & Holbrook, 2012). However, a different effect was observed on a measure of the time taken to eventually *go past* the critical region (Regression Path). This measure of reading time captures early processing plus any initial "repair" processes where the region is reread after an initial regressive saccade. Regression Path times revealed that the effect of desire was quickly moderated by the protagonist's beliefs. There was an effect of desire when the protagonist *knew* how to achieve her desires (i.e., longer reading times when outcome was desired vs. not desired), but no such effect when the protagonist did not know how to achieve her desire.

Effects emerging later in the text (First Pass Regressions Out of the post-critical region and Regressions In to the critical region) were influenced only by the protagonist's desire state. Total reading time (which combines both early and late processing) revealed an interaction similar to that observed on Regression Path times, but this interaction was more symmetrical due to effects of desire that emerged later in the text.

The temporal pattern of effects is consistent with the idea that desires are initially at the forefront of the readers mind. Inferences based solely on desire are then quickly moderated by what the reader knows about the protagonist's beliefs, but as time goes on this inhibition dissipates, with desires returning to the forefront. This suggests that the inhibition of privileged desire-based knowledge may be an effortful process that is not systematically sustained in later steps of processing. Early and late processes were consistent with evidence that desires are more salient in the minds of observers than beliefs (Malle & Holbrook, 2012) resulting in interference that has previously been described as the *curse of knowledge* (Birch & Bloom, 2007). Intermediate processes were consistent with an account in which readers only infer intent when the protagonist has both the relevant desires and beliefs, but not when one of these states is absent (cf. Malle & Knobe, 1997).

One alternative explanation for the strong effect of desire is that in our vignettes the sentence describing beliefs (sentence 2) was always presented before the sentence describing desires (sentence 3), which in turn was presented immediately before the critical sentence (sentence 4). This sequence might cause desires to be foregrounded in the reader's mind just before the critical sentence is encountered, whereas beliefs may be backgrounded to some extent by this intervening text. However, this type of backgrounding is typically only found when there are multiple sentences of intervening text (e.g., Albrecht & Myers, 1995; Lea, Mulligan, & Walton, 2005). One way to test this alternative account would be to reverse or counterbalance the order in which beliefs and desires are presented in our vignettes (e.g., desires presented in sentence two and beliefs presented in sentence 3). Crucially, such a possibility does not detract from our conclusion that the saliency of desires can interfere with what readers know about a protagonist's belief state (regardless of whether the generally greater salience of desires described by Malle and Holbrook may have been aided to some extent by the design of our materials).

Our findings have implications for researchers studying narrative comprehension, theory of mind, and conditional reasoning. Narrative comprehension studies have shown that readers can track the beliefs (de Vega et al., 1997), desires (Huitema et al., 1993; Poynor & Morris, 2003), and intentions (Zwaan et al., 1995) of a protagonist. We build on these findings by independently manipulating both beliefs and desires to reveal how and when readers use these sources of information to infer the intentions of a protagonist. At the most general level of description, readers in our experiment detected an inconsistency between their expectations and the behavior of the protagonist. In that sense, the cognitive processes at work in our task are presumably at work in other cases of semantic or pragmatic inconsistency. The neural signatures of inconsistency detection must be detectable in our task, just as they are for other types of inconsistency. The social specificity of our results is in the input of this inconsistency detection process, as it would be impossible to detect an inconsistency if ToM had not delivered first an expectation about the intention of the protagonist.

Our results are also relevant to studies of conditional reasoning, particularly those that focus on the role that utility can play when inferring behavior (e.g., Bonnefon, 2009). Future studies would benefit from considering not just what the protagonist in a reasoning problem knows, but also what the participant knows, as both factors may influence the conclusions that participants are willing to endorse at different temporal stages of the reasoning process. Our findings are also relevant to more traditional methods of studying conditional inference. Traditional conditional inference tasks provide a major conditional premise ("if Alice renewed over the Internet she would save £100") a minor premise (Alice renewed over the Internet) and a conclusion or outcome (she saved £100). Our materials essentially provided a major premise and a minor premise, with the conclusion being implicit. Previous research has shown that when a minor premise is read the conclusion is spontaneously inferred (Lea, 1995), but future research would benefit from presenting the minor premise and conclusion separately so that effects can be more accurately localized.

A promising avenue for future research is to investigate the moderators and boundary conditions of this result, based on the moderators and boundary conditions of belief-desire inferences identified in reasoning research. For example, we know that reasoners are influenced by the causal structure of a main premise such as "if Alice renews her insurance over the internet, she benefits from a good service" (Bonnefon & Sloman, 2013). The direction of causality (does renewing over the Internet result in better service, or is renewing over the Internet a reason to believe that the service is good?), the existence of disablers (reasons why Alice might not get a good service even after renewing online) or alternatives (ways to get a good service without renewing online), can all influence the reasoners' willingness to infer that Alice will renew her insurance over the Internet (Corner, Hahn, & Oaksford, 2011; Evans et al., 2008). Specifically, these factors can all reduce the likelihood that readers will infer an intention. If causality works

backwards (i.e., renewing over the internet is a consequence of having a good policy), reasoners no longer infer that Alice will renew online; if renewing online does not guarantee a good service, reasoners no longer infer that Alice will renew online; and if there are other options for getting good services, reasoners are no longer sure that Alice will choose the option of renewing online. An important question is whether these factors influence fast ToM inferences to the same extent that they influence conclusions in slow-paced reasoning experiments. We look forward to this fuller integration of research on ToM, reasoning, and narrative comprehension.

From this study we can conclude that inferring intent from the beliefs and desires of others is a complex operation. Experience tells us that adults are generally adept at predicting the behavior of others, but a growing number of experimental studies have shown that intent inferences are susceptible to interference from privileged knowledge. Our data suggest that this interference occurs very rapidly, but is moderated during intermediate stages of processing. However, this moderation is not systematically sustained, suggesting that it might be an effortful process. Failure to maintain to this effortful suppression of desire-based inferences may be one reason why people are susceptible to the *curse of knowledge*.

Acknowledgments

This work was supported by an Experimental Psychology Society Small Grant awarded to the first author. We thank Jamal Kinsella for collecting the eye movement data and Andrew J. Stewart for comments on a previous draft of this manuscript.

Electronic Supplementary Material

The electronic supplementary material is available with the online version of the article at http://dx.doi.org/10.1027/1618-3169/a000290

ESM 1. Appendix.
Experimental material used in this study.
ESM 2. Excel lists.
Raw data of all means by subjects and items.

References

Albrecht, J. E., & Myers, J. L. (1995). Role of context in accessing distant information during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 1459–1468.

Apperly, I. A. (2012). What is "theory of mind"? Concepts, cognitive processes and individual differences. *The Quarterly Journal of Experimental Psychology, 65*, 825–839.

Birch, S. A. J., & Bloom, P. (2007). The curse of knowledge in reasoning about false beliefs. *Psychological Science, 18*, 382–386.

Bonnefon, J. F. (2009). A theory of utility conditionals: Paralogical reasoning from decision-theoretic leakage. *Psychological Review, 116*, 888–907.

Bonnefon, J. F. (2012). Utility conditionals as consequential arguments: A random sampling experiment. *Thinking and Reasoning, 18*, 379–393.

Bonnefon, J. F., Girotto, V., & Legrenzi, P. (2012). The psychology of reasoning about preferences and unconsequential decisions. *Synthese, 187*, 27–41.

Bonnefon, J. F., Haigh, M., & Stewart, A. J. (2013). Utility templates for the interpretation of conditional statements. *Journal of Memory and Language, 68*, 350–361.

Bonnefon, J. F., & Hilton, D. J. (2004). Consequential conditionals: Invited and suppressed inferences from valued outcomes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30*, 28–37.

Bonnefon, J. F., & Sloman, S. A. (2013). The causal structure of utility conditionals. *Cognitive Science, 37*, 193–209.

Corner, A., Hahn, U., & Oaksford, M. (2011). The psychological mechanism of the slippery slope argument. *Journal of Memory and Language, 64*, 153–170.

de Vega, M., Díaz, J. M., & León, I. (1997). To know or not to know: Comprehending protagonists' beliefs and their emotional consequences. *Discourse Processes, 23*, 169–192.

Egan, S. M., & Byrne, R. M. J. (2012). Inferences from counterfactual threats and promises. *Experimental Psychology, 59*, 227–235. doi: 10.1027/1618-3169/a000147

Evans, J. S. B. T., Neilens, H., Handley, S. J., & Over, D. E. (2008). When can we say "if"? *Cognition, 108*, 100–116.

Ferguson, H. J., & Breheny, R. (2011). Eye movements reveal the time-course of anticipating behaviour based on complex, conflicting desires. *Cognition, 119*, 179–196.

Ferguson, H. J., & Breheny, R. (2012). Listeners' eyes reveal spontaneous sensitivity to others' perspectives. *Journal of Experimental Social Psychology, 48*, 257–263.

Huitema, J. S., Dopkins, S., Klin, C. M., & Myers, J. L. (1993). Connecting goals and actions during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 1053–1060.

Keysar, B. (1994). The illusory transparency of intention: Linguistic perspective taking in text. *Cognitive Psychology, 26*, 165–208.

Keysar, B., Barr, D. J., Balin, J. A., & Brauner, J. S. (2000). Taking perspective in conversation: The role of mutual knowledge in comprehension. *Psychological Science, 11*, 32–38.

Kovács, A., Téglás, M., & Endress, A. D. (2010). The social sense: Susceptibility to others' beliefs in human infants and adults. *Science, 330*, 1830–1834.

Lea, R. B. (1995). On-line evidence for elaborative logical inferences in text. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 1469–1482.

Lea, R. B., Mulligan, E. J., & Walton, J. L. (2005). Accessing distant premise information: How memory feeds reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 387–395.

Malle, B. F. (1999). How people explain behavior: A new theoretical framework. *Personality and Social Psychology Review, 3*, 23–48.

Malle, B. F., & Holbrook, J. (2012). Is there a hierarchy of social inferences? The likelihood and speed of inferring intentionality, mind, and personality. *Journal of Personality and Social Psychology, 102*, 661–684.

Malle, B. F., & Knobe, J. (1997). The Folk Concept of Intentionality. *Journal of Experimental Social Psychology, 33*, 101–121.

Miller, D. T. (1999). The norm of self-interest. *American Psychologist, 54*, 1053–1060.

Pickering, M. J., Frisson, S., McElree, B., & Traxler, M. (2004). Eye movements and semantic composition. In M. Carreiras & C. Clifton (Eds.), *The on-line study of sentence comprehension: Eyetracking, ERP, and beyond* (pp. 33–50). New York, NY & London, UK: Psychology Press.

Poynor, D. V., & Morris, R. K. (2003). Inferred goals in narratives: Evidence from self-paced reading, recall, and eye movements. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*, 3–9.

Qureshi, A., Apperly, I. A., & Samson, D. (2010). Executive function is necessary for perspective selection, not level-1 visual perspective calculation: Evidence from a dual-task study of adults. *Cognition, 117*, 230–236.

Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin, 124*, 372–422.

Samson, D., Apperly, I. A., Braithwaite, J. J., Andrews, B. J., & Bodley Scott, S. E. (2010). Seeing it their way: What other people see is calculated by low-level and early acting processes. *Journal of Experimental Psychology: Human Perception and Performance, 36*, 1255–1266.

Schneider, D., Bayliss, A. P., Becker, S. I., & Dux, P. E. (2012). Eye movements reveal sustained implicit processing of others' mental states. *Journal of Experimental Psychology: General, 141*, 433–438.

van der Wel, R. P., Sebanz, N., & Knoblich, G. (2014). Do people automatically track others' beliefs? Evidence from a continuous measure. *Cognition, 130*, 128–133.

Wellman, H. M., & Woolley, J. D. (1990). From simple desires to ordinary beliefs: The early development of everyday psychology. *Cognition, 35*, 245–275.

Zwaan, R. A., Langston, M. C., & Graesser, A. C. (1995). The Construction of Situation Models in Narrative Comprehension: An Event-Indexing Model. *Psychological Science, 6*, 292–297.

Received July 1, 2014
Revision received November 20, 2014
Accepted November 21, 2014
Published online May 7, 2015

Matthew Haigh

Department of Psychology
Northumbria University
Newcastle upon Tyne
NE1 8ST
UK
Tel. +44 191 227-3472
E-mail matthew.haigh@northumbria.ac.uk

