Article

Personality and Social Psychology Bulletin
1–15
© 2017 by the Society for Personality and Social Psychology, Inc
Reprints and permissions: sagepub.com/journalsPermissions.nav
DOI: 10.1177/0146167217702375
pspb.sagepub.com
SAGE

Information-Acquisition Processes in Moral Judgments of Blame

Steve Guglielmo[1] and Bertram F. Malle[2]

Abstract

When people make moral judgments, what information do they look for? Despite its theoretical and practical implications, this question has largely been neglected by prior literature. The recent Path Model of Blame predicts a canonical order in which people acquire information when judging blame. Upon discovering a negative event, perceivers consider information about causality, then intentionality, then (if the event is intentional) reasons or (if the event is unintentional) preventability. Three studies, using two novel paradigms, assessed and found support for these predictions: In constrained (Study 1) and open-ended (Study 2) information-acquisition contexts, participants were most likely, and fastest, to seek information in the canonical order, even when under time pressure (Study 3). These findings indicate that blame relies on a set of information components that are processed in a systematic order. Implications for moral judgment models are discussed, as are potential roles of emotion and motivated reasoning in information acquisition.

Keywords

blame, morality, information seeking, intentionality, mental states

Received September 27, 2016; revision accepted March 5, 2017

What information do people seek when making moral judgments, and in what order? In legal and media contexts, people often have minimal or incomplete information and must discover new information that guides their eventual judgments of blame and punishment. Beyond such practical implications, assessing the process of information acquisition is of theoretical importance, because it helps to clarify the conceptual and psychological structure that underlies people's moral judgments. Surprisingly, however, information-acquisition processes in moral judgment have been overlooked in the existing literature.

Moral judgments, like any other judgments, rely on information processing. An information-processing framework of moral judgment would specify the *information input* that guides people's judgments and the *psychological processes* that operate on this information to generate the judgments (Guglielmo, 2015). The current studies elucidate this information-processing basis of blame by investigating what information people acquire, and in what order, en route to blame judgments.

The Role of Information in Moral Judgment

Many early models of moral judgment focused on which types of information are important for perceivers' judgments. Shaver (1985) argued that responsibility judgments are guided by causality, wrongness, general controllability (that the event is something that agents generally can control), and specific intentionality (that this specific behavior was intentional). Weiner (1995) focused on controllability, suggesting that "the cause must be controllable if the person is to be held responsible" (Weiner, 1995, p. 11). Schlenker, Britt, Pennington, Murphy, and Doherty (1994) likewise emphasized controllability, while also suggesting that responsibility judgments reflect the agent's obligation and the overall wrongness of the event.

Judgments of responsibility have proven to be ambiguous—referring variably to causality, role-based obligations, or simply blame—and existing models of responsibility face additional concerns. Weiner (1995) and Schlenker et al. (1994) both omit intentionality; the concept of controllability typically applies to both intentional and unintentional behaviors, even though people's moral judgments robustly differentiate between these two (Cushman, 2008; Young & Saxe, 2009). Moreover, these early models have little to say about the processes involved in moral judgment. Classic vignette methodologies, in which all information is presented unambiguously and simultaneously, focus on the final product of

[1]Macalester College, Saint Paul, MN, USA
[2]Brown University, Providence, RI, USA

Corresponding Author:
Steve Guglielmo, Macalester College, 1600 Grand Ave., Saint Paul, MN 55105, USA.
Email: sgugliel@macalester.edu

information processing—the judgment—but not the dynamic, complex nature of the information processing itself.

Cushman's (2008) more recent model shifted away from the ambiguous judgments of responsibility, instead focusing on blame and wrongness, while further clarifying their underlying information features. Blame is a particularly rich moral judgment, requiring greater and more differentiated information processing than other judgments such as responsibility or wrongness (Malle, Guglielmo, & Monroe, 2014). Cushman's (2008) model posits that blame judgments are driven by considerations of *consequences* (what actually happened as a result of an agent's action) and *mental states* (what the agent believed or wanted to happen). This model is consistent with evidence demonstrating the robust influence of mental states on blame among adults (Gray, Young, & Waytz, 2012; Lagnado & Channon, 2008; Young & Saxe, 2009) and children (Nelson, 1980; Zelazo, Helwig, & Lau, 1996). It also, unlike several prior models, incorporates the fact that even unintentional negative behaviors can incur blame.

Other theoretical models have focused less on the types of information that inform moral judgments and instead on the psychological processes that generate those judgments. Some models argue that moral judgments are driven largely by intuitive processing (Alicke, 2000; Haidt, 2001), whereas others assert that both intuitive and deliberative processing together shape moral judgments (Greene, 2007). Furthermore, a host of evidence has shown that emotion can influence moral judgments, as revealed by studies that induce emotion (Horberg, Oveis, Keltner, & Cohen, 2009) or that examine patients with emotional deficits (Moretto, Làdavas, Mattioli, & di Pellegrino, 2010). However, these investigations have focused primarily on the distinction between cognition and affect and have been less concerned with how information processing (whether "cognitive" or "affective") integrates the kinds of information (e.g., intentionality, mental states) that guide moral judgments.

An information-processing perspective on moral judgment raises many questions that have not been answered yet—from the acquisition of information to its integration and use in judgment. We focus here on the information-acquisition phase and investigate these questions about judgments of blame: What are the information components that underlie blame judgments, and how do people acquire this information? Are the components processed in any order available, or is there a more systematic, efficient sequence of processing? To investigate these questions, we need both a theoretical model that provides clear predictions about information acquisition and a methodology that allows perceivers to engage in a dynamic search for information.

Sequential Information Acquisition

The Path Model of Blame (Malle et al., 2014; Malle, Guglielmo, & Monroe, 2012) offers predictions about the process of information acquisition, because it explicitly specifies the canonical information-processing sequence that guides blame. The model, depicted in Figure 1, asserts that information processing toward blame begins with a social perceiver detecting a norm-violating event. The perceiver then assesses causality, determining who or what caused the event. If the event appears agent-caused, the perceiver determines whether it was intentional or not. If intentional, the degree of blame depends on the agent's reasons for acting; if unintentional, the degree of blame depends on whether the agent could have *prevented* it.[1]

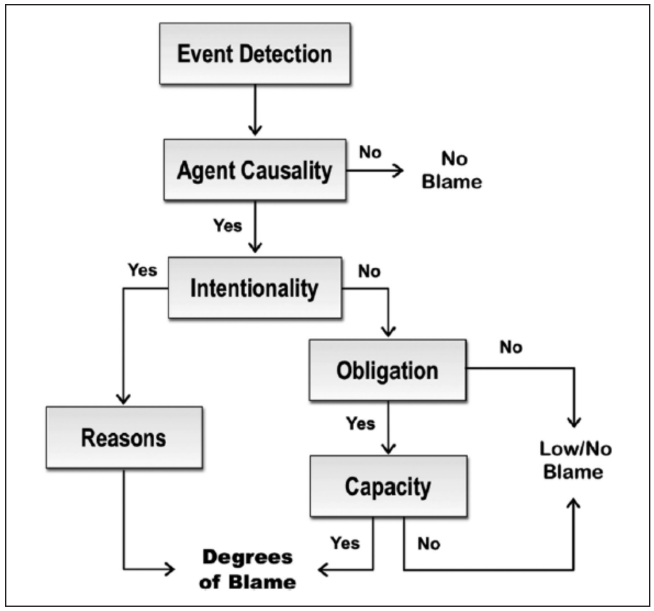

Figure 1. Path Model of Blame (Malle, Guglielmo, & Monroe, 2014).

Much prior research supports the importance of these individual components: blame is greater when an agent is the cause of a negative outcome (Cushman, 2008; Sloman, Fernbach, & Ewing, 2009), acts intentionality (Lagnado & Channon, 2008; Young & Saxe, 2009), possesses undesirable reasons or motives (Inbar, Pizarro, & Cushman, 2012; Woolfolk, Doris, & Darley, 2006), and has the capacity to prevent the outcome (Alicke, Weigold, & Rogers, 1990; Weiner, 1995). The Path Model predicts that across complementary methodological contexts, people will not only seek information about these components generally but will do so in a particular order.

Competing accounts dispute the plausibility of such sequential processing of blame. For example, Alicke (2014) argues that "a problem with [sequential] models is the assumption that blame is ordinarily evaluated in such discrete stages" (p. 190). Similarly, Schein and Gray (2014) suggest that "moral cognition cannot be confined to discrete boxes, whether in structure or in sequence" (p. 236). According to these accounts of "nonsequential" processing—and in contrast to the predictions of the Path Model—there is little reason to expect any specific sequence or order of information processing en route to blame. This is an empirically testable dispute, so we need an experimental paradigm that enables such a test.

A Method to Study Information Acquisition

Information-acquisition strategies have been studied extensively in the broader literature on judgment and decision making (Dhami & Harries, 2010; Gigerenzer & Gaissmaier, 2011). One method of assessing such strategies is an "information display board," in which participants can search through a matrix of information, revealing whatever information will help them complete their task (Jacoby, Jaccard, Kuss, Troutman, & Mazursky, 1987). For example, participants might be able to view particular kinds of information about various apartments (price, size, parking arrangements, etc.) to help them select their preferred one (Payne, 1976). Another method is an "active information search" paradigm, in which participants ask questions and receive answers, thus providing a more naturalistic method for uncovering the features most critical for their decision strategies (Reisen, Hoffrage, & Mast, 2008). The current studies implement variants of both methodologies to shed light on the information-acquisition process of blame. Importantly, these methods, rather than surveying people about their implicit theories of information acquisition, allow people actually to search for information and update their knowledge in service of arriving at a judgment.

Overview and Predictions

The Path Model of Blame (Malle et al., 2014) makes three kinds of predictions about information acquisition toward blame: what specific information components people will focus on, in what canonical order, and with what structural dependencies. In brief, when people discover a norm-violating event, they will focus on agent causality information, then intentionality information, then reasons information (if the event was considered intentional) or preventability information (if the event was considered unintentional). In contrast, alternative models of moral judgment either make no specific predictions about the components of information processing toward blame (e.g., Haidt, 2001; Gray et al., 2012), make different predictions about the components (e.g., Weiner, 1995 focuses on controllability and responsibility judgments) or suggest that perceivers are not bound by any particular processing order when determining blame (Alicke, 2000, 2014; Schlenker et al., 1994).

We test these predictions in three studies, each providing a different information-acquisition context. In the "information selection" paradigm of Study 1, participants decide whether to accept or reject offers of various types of information. This paradigm constrains the types of information that people can acquire and tests competing predictions about the processing order and mutual dependencies among these types. In the "information search" paradigm of Study 2, participants actively request whichever information they want for assigning blame. This paradigm thus tests competing predictions both about the specific types of information that people care about and their processing order and dependencies. Last, Study 3 incorporates a time pressure manipulation into the information selection paradigm, assessing whether systematic information requires deliberative processing.

The combined sample sizes of the studies reported here (total $n = 172$) enabled us to detect within-subject effect sizes of $d = .20$ ($\alpha = .05$) with $\beta = .74$ (Faul, Erdfelder, Lang, & Buchner, 2007).

Study 1

In this study, participants read about partially described negative events and then respond to offers for expanded information about each event. Different categories of initial events "place" participants at particular points of the hypothesized information-processing order en route to blame (see Figure 1): The key categories are *Outcome* events (describing only a negative outcome), *Causality* events (describing a person-caused negative outcome), *Intentional* events (describing an intentional negative action), and *Unintentional* events (describing an unintentional negative behavior). According to the Path Model, each of these events demands specific kinds of subsequent information components most useful in forming a blame judgment. For example, for an Outcome event, people should be most interested in information about causality, because it is the next information component activated after the detection of a negative event; for a Causality event, people should be most interested in information about intentionality. The Path Model thus predicts that people should be faster and more likely to accept offers about the information component that comes next in the sequence postulated by the model ("Next" offers), compared with information components that skip beyond the next component ("Jump" offers) or that revert to a prior component ("Back" offers). Models that are agnostic about sequential information processing (Cushman, 2008), or that explicitly deny such sequencing (Alicke, 2014; Schein & Gray, 2014), predict no such patterns of differential likelihood and speed of accepting information.

Method

Participants and procedure. U.S. participants ($n = 61$) were recruited from Amazon Mechanical Turk and received $.50 for completing the study online. They learned that they would read a series of events, that their task was to "determine whether someone deserves blame (and how much) for what happened," and that they would have opportunities to acquire more information. They then completed three practice trials and 32 experimental trials, separated by three breaks, each of which provided a brief reminder of the task instructions. See Supplementary Material Appendix A for the verbatim task instructions used in the studies reported here.

Each trial consisted of four steps. Once the participant pressed the Space Bar to begin the trial, the initial event description appeared (Step 1). After 2,000 ms, an "information offer" signal appeared, signifying that an offer was coming (the event description remained on the screen); once the participant pressed the space bar, the actual information offer appeared for a maximum of 6,000 ms (Step 2). Participants responded by pressing Yes (accepting the offer: Letter I) or No (rejecting the offer: Letter O), and response times were recorded. Step 3 displayed the relevant information (if the offer was accepted) or removed the offer from the screen (if it was rejected). Each trial contained between one and three information offers, depending on the initial event type (e.g., Causality events could have maximally three offers; Intentional events, maximally one).

After the final offer of the trial, participants assigned blame (Step 4) by answering a yes/no question and, if they responded yes, providing a 1 to 9 rating.

Materials. Each trial presented an event description designed to put participants into a particular information state, corresponding to a "place" within the information-processing path to blame: Outcome (describing a negative outcome), Causality (describing an agent-caused negative outcome), Intentional (intentional negative behavior), Unintentional (unintentional negative behavior), *Reasons* (intentional behavior with information about the agent's reasons), *Preventability* (unintentional behavior with information about whether the agent could have prevented it), and *Nonagent* (negative outcome occurring from natural causes). See Supplementary Material Appendix B for all event descriptions.

Each offer of information was represented by one-word *information cues*: CAUSE (information about who or what caused the event), INTENTIONAL (information about whether the event was intentional or unintentional), REASONS (information about a person's reasons for acting), and PREVENTABLE (information about whether someone could have prevented the event). For example, for the Outcome event "Wayne had a black eye," the CAUSE information was "Wayne had a black eye, which was somehow caused by Charles," and the INTENTIONAL information was "Wayne had a black eye, which Charles intentionally gave him."

To facilitate analysis, information offers were categorized into three theoretically relevant *offer types*, for which the Path Model makes specific predictions:

- *Next* offers provided information about the subsequent component in the path to blame. For example, for an Outcome event, a CAUSE cue constituted a Next offer; for a Causality event, an INTENTIONAL cue constituted a Next offer.[2]

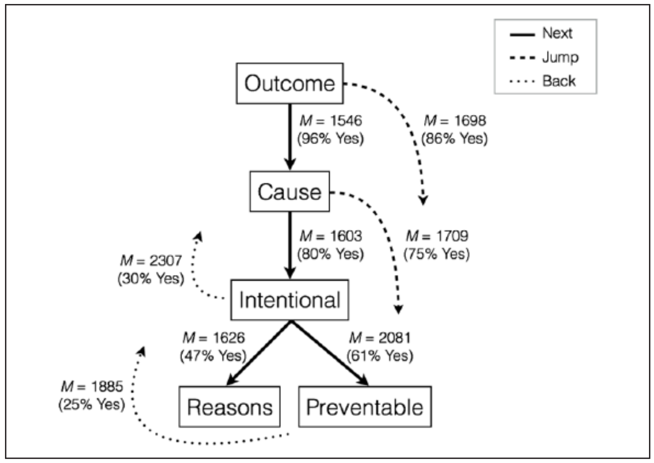

Figure 2. Average acceptance response times (and acceptance response rates) from each information state as a function of offer type (Study 1).

- *Back* offers provided information about a component further back in the path. For example, for an Intentional event, a CAUSE cue constituted a Back offer.
- *Jump* offers provided information about a component two or more steps down the path. For example, for an Outcome event, an INTENTIONAL, REASONS, or PREVENTABLE cue would all constitute Jump offers.

The Path Model predicts that people will be more likely, and faster, to accept information that targets the next component of the model (Next offers), as compared with other offer types.

Figure 2 presents a schematic of the possible offer types at each information state. At the Outcome state, Next offers (for Cause information) and Jump offers (for Intentional, Reasons, or Preventable) were possible. Likewise, at the *Cause* state, Next offers (for Intentional) and Jump offers (for Reasons or Preventable) were possible. At the Intentional and Unintentional states, Next offers (for Reasons and Preventable, respectively) and Back offers (for Cause) were possible. Finally, at the Reasons and Preventability states, only Back offers (for Cause or Intentional) were possible.

Results

Data preparation. All trials (N = 3,036 trials with yes or no responses, across 61 participants) were screened for response time outliers via a two-stage procedure. First, responses faster than 400 ms were excluded (1.2% of all trials). Second, within-subject outliers were excluded (responses exceeding within-subject mean ± 2.5 within-subject *SD*s; 2.8% of all trials). This process yielded 2,914 valid trials for further analysis.

Next versus Back. We first assessed the prediction that people accept offers moving them down the information-processing path to blame (i.e., Next offers) more often than those moving them back up the path (i.e., Back offers). A binomial multilevel model was specified with trials as the unit of analysis and response (yes vs. no) modeled as a function of offer type (Next vs. Back), information cue (Cause vs. Intentional: the two cues for which Next and Back offers were both possible), and the Offer Type × Information Cue interaction (all Level 1 predictors), including random intercepts for subjects (at Level 2). People accepted Next offers far more often (88% of such offers were accepted) than Back offers (26% were accepted), $z = 17.24$, $p < .001$, 95% confidence interval (CI) = [.562, .669], Cohen's $h = 1.36$. There was also an Offer Type × Information Cue interaction, whereby the Next preference was stronger for causality information (96% vs. 22%) than for intentionality (80% vs. 33%), $z = 6.25$, $p < .001$.

A similar multilevel model examined acceptance response times (i.e., trials in which participants accepted the offered information), again including offer type, information cue, and their interaction as Level 1 predictors, and random intercepts for subjects at Level 2. People accepted Next offers more quickly ($M = 1,572$ ms, $SD = 827$ ms) than Back offers ($M = 1,987$, $SD = 1,100$), $t(841) = 4.52$, $p < .001$, 95% CI = [293.8, 744.0], $d = .49$, and this effect was not moderated by information cue.

Next vs. Jump. The Path Model predicts that people will accept Next offers more often than Jump offers. For example, people learning about a negative outcome should accept a Cause offer (Next) more often than an Intentional offer (Jump). The design allowed two tests of this hypothesis, each at a different information state along the path to blame: (a) at Outcome, Cause offer (Next) versus offers that jump beyond Cause, and (b) at *Cause*, Intentional offer (Next) versus offers that jump beyond Intentional. Binomial multilevel models examined trial-level responses as a function of offer cue, including random intercepts for subjects. As shown in Figure 2, at the Outcome state, people accepted Next offers (for causality information) more often (96%) than Jump offers (86%), $z = 4.82$, $p = .01$, 95% CI = [.056, .131], $h = 0.33$. At the *Cause* state, people accepted Next offers (for intentionality information) slightly more often (80%) than Jump offers (75%), $z = 1.72$, $p < .09$, 95% CI = [−.018, .113], $h = 0.13$.

Similar analyses assessed response time patterns and suggest that it takes additional time to infer jumped-over information (see Figure 2). At the Outcome state, people accepted Next offers (for causality information) more quickly ($M = 1,546$, $SD = 779$) than Jump offers ($M = 1,698$, $SD = 824$), $t(752) = 3.32$, $p < .001$, 95% CI = [60.6, 234.9], $d = .19$. At the *Cause* state, people accepted Next offers (for intentionality information) slightly more quickly ($M = 1,603$, $SD = 881$) than Jump offers ($M = 1,709$, $SD = 825$), $t(482) = 1.84$, $p < .07$, 95% CI = [−8.6, 255.5], $d = .12$.

One unexpected pattern emerged regarding preventability information in particular. People were slower to accept preventability information when it came as a Next offer (from the Unintentional state: $M = 2,081$, $SD = 931$) than as a Jump offer (from either the Outcome or *Cause* state: $M = 1,568$, $SD = 700$), $t(194) = 5.29$, $p < .001$, 95% CI = [308.7, 675.4], $d = .61$.

Discussion

Study 1 allowed participants to accept or reject information components widely considered to be relevant for blame: causality, intentionality, reasons, and preventability. In contrast to models that either doubt or make no claims about sequential processing, the Path Model holds that people seek these components in a canonical order, thus predicting that people will more often and more quickly seek the next component down the path (i.e., Next offers) as compared with components further up (Back offers) or further down the path (Jump offers).

The results supported these hypotheses. People accepted Next offers far more often than Back offers, and when they did accept Back offers, it took them substantially longer. People also tended to accept Next offers more often and more quickly than Jump offers, though these patterns were somewhat less robust. One reason for these weaker effects in the Next versus Jump comparison is that Jump offers may be seen as implying relevant "Next" information. If one can trust the information source, then one might infer from an offer to learn about an agent's reasons that the agent indeed acted intentionally. Consistent with this interpretation, the reaction times for Jump offers were slower than for Next offers, hinting at the added inference that participants may have made.

The results demonstrate that the constitutive information components of blame are not on equal footing. People did not simply accept any type of information whenever they could get it. Rather, they showed clear preferences in their information acquisition: causality information had processing priority over intentionality information, which typically had processing priority over reasons and preventability information. These patterns are consistent with the Path Model but not with nonsequential accounts.

One limitation of the information selection paradigm of Study 1 is that participants responded to information offers that were framed in terms of a priori information components. These components have been postulated, in various subsets, by multiple models, and as a set they are the ones specified by the Path Model of Blame. However, this paradigm does not allow us to assess whether these are indeed the information components people really care about when forming judgments of blame. Other components, such as responsibility (Schlenker et al., 1994), controllability (Weiner, 1995), or character (Pizarro & Tannenbaum, 2011), may be important to people, but they had no opportunity to select them as potential information content. By allowing participants to engage in an unconstrained search for information, we can test these competing predictions about central information categories and provide another, even more robust, test of the predicted processing order and dependencies.

Study 2

Study 2 implements an open-ended "information search" paradigm in which participants can ask about any information they want en route to assigning blame. The Path Model predicts a specific set of information types people seek out (norm-violating event, agent causality, intentionality, reasons, preventability), a canonical ordering, and a bifurcation of information depending on intentionality (reasons information for intentional events, preventability information for unintentional events). Nonsequential accounts make no such predictions about ordering or bifurcated processing.

Method

Participants and procedure. Undergraduate participants (n = 32) completed the study individually for course credit. They were told that their task for each event was to "assign blame (and how much) for what happened" and that they could "learn more by asking questions. These must be posed as Yes/No questions and the experimenter will look up their correct answers in a database that contains all the information about each event." To avoid leading them to ask about specific concepts, participants were not given any example questions.

The experimenter and the participant were seated in separate rooms; the participant then completed 14 trials, each consisting of four steps. Via a computer chat program, the experimenter sent the initial event description to the participant (Step 1), followed by a scripted prompt instructing the participant to ask a question or reply "ready" if ready to assign blame (Step 2). The participant asked a question verbally or said "ready" (Step 3). If the participant asked a question, the experimenter sent a scripted prompt indicating that the answer was being retrieved, in fact retrieved the answer from the database, and then sent the yes/no answer after 6 to 8 s.[3] Participants could ask up to four questions about each event. Thus, Steps 2 and 3 (question prompt and participant response) were repeated until the participant said "ready" or had asked four questions, at which point the participant rated blame on a 0 to 10 scale (Step 4).

Design and materials. In each trial, participants received a specific category of starting event descriptions: Outcome events (describing a negative outcome), Causality events (describing an agent-caused negative outcome), or *Intentionality* events (describing an Intentional or Unintentional negative behavior). Each event in the Outcome and Causality category was assigned to be "ultimately" intentional or unintentional—that is, the database that contained the answers to participants' questions designated the event, a priori, as either intentional or unintentional. Outcome events could also be ultimately nonagent caused. Thus, participants encountered seven different types of event descriptions, each implemented in two trials (for a total of 14 trials): Outcome (ultimately intentional, unintentional, or nonagent), Causality (ultimately intentional or unintentional), *Intentional*, and *Unintentional*.

To ensure that search strategies for informationally sparse events (e.g., Outcome events) would not be biased by participants' processing of informationally rich events (e.g., Intentionality events), event descriptions were presented in blocks of increasing informational detail: first Outcome events, then Causality events, then Intentionality events. The order of ultimately intentional versus unintentional events was counterbalanced across participants.

Data preparation and coding. Based on prior literature and bottom-up examination of the questions, 24 initial categories were derived (Supplementary Material Appendix C), into which all questions (N = 1,124 across 32 participants) could be coded.

Two independent coders, one of whom was blind to all hypotheses, coded 119 questions (intercoder agreement = 86%); after resolving disagreements, they coded another 121 questions (intercoder agreement = 79%). Each coder then coded approximately half the remaining questions; Coder 1 then independently coded 399 of Coder 2's questions (intercoder agreement = 79%) and Coder 2 independently coded 125 of Coder 1's questions (intercoder agreement = 85%). Disagreements were again jointly resolved. Finally, inspection of the categories warranted the inclusion of two additional categories (*event clarification* and *uncodable*), for a total of 26 categories. The coders independently reconsidered all questions for possible reclassification into these new categories.

The 26 surface-level categories were then collapsed into nine conceptually distinct umbrella categories: event clarification, negativity, causality, intentionality, reasons, oblication, preventability, response, and uncodable (Table 1 and Supplementary Material Appendix D). For example, surface-level questions about *carelessness, recklessness*, and *incapacitating factors* constituted the umbrella category of preventability.[4]

Results

Categories of information search. Several previous models of moral judgment suggest that people might search for information about controllability (Weiner, 1995), character or prior negative behavior (Alicke, 2000), and responsibility (Schlenker et al., 1994; Shaver, 1985; Weiner, 1995). Participants, however, rarely asked about these features. Just four out of 1,124 questions (under 1%) asked about "control," seven (under 1%) asked about prior negative behavior or general behavioral tendencies (surface-level code *occur_agent*), and 18 (1.6%) asked about "responsibility" (or the variant "responsible").[5]

The vast majority of questions targeted components that, at least as subsets, appear in some models of blame: causality (25% of all questions), intentionality (21%), event clarification

Table 1. Definitions and Examples of Umbrella Coding Categories in Study 2.

Category	Definition	Examples
Event clarification	Contextual or background information (not clearly targeting any other category)	1. Were they in some sort of physical fight? 2. Did the other person use a weapon to hurt Monica?
Negativity	Scope of consequences or presence of other negative consequences.	1. Did anyone get sick from eating food at the restaurant? 2. Did Brittany survive the knife wound in her leg?
Causality	Who/what caused event (including person's characteristics, background, etc.).	1. Was the forest fire caused by a person? 2. Did he die of natural causes?
Intentionality	Whether event was caused intentionally (purposely, deliberately, accidentally, etc.).	1. Did Bill shatter the window deliberately? 2. Did Ryan cause the house to burn down on purpose?
Reasons	Agent's reasons for bringing about the event (including whether agent was justified).	1. Was it done out of animosity? 2. Was Paul trying to help Brittany?
Preventability	Whether agent could have prevented event (including carelessness, recklessness).	1. Could Wayne have prevented that accident? 2. Was Ryan being careless when he accidentally started the fire?

(16%), reasons (14%), and preventability (12%). Together, nearly 90% of all questions targeted these components, which collectively correspond closely to the conceptual framework of the Path Model of Blame.

Depth of information search. A first hint at a canonical processing sequence lies in the fact that participants asked the greatest number of questions about Outcome events (M = 2.80 questions per event), followed by Causality events (M = 2.39), Unintentional events (M = 1.62), and Intentional events (M = 1.48). All subsequent analyses examine trials in which at least one question was asked (true for 92% of Outcome, 88% of Causality, 64% of Unintentional, and 56% of Intentional events) and treat individual questions (N = 1,124) as the unit of analysis.

Bifurcation of intentionality. The Path Model holds that some information components (causality, intentionality) are relevant for any negative event, whereas others (reasons, preventability) are differentially relevant depending on whether the event was intentional or not. We can assess these predictions by examining people's questions about "ultimately" intentional versus unintentional events. As the Path Model predicts, questions about early stages of the model (event clarification, other negativity, causality, intentionality) were equally prevalent across ultimate intentionality, all $\chi^2(1) <$ 1.9, ps > .10. Together, these four question types constituted 61% of all questions for ultimately intentional events and 67% of all questions for ultimately unintentional events. Other question types, however, were differentially prevalent depending on ultimate intentionality. Questions about reasons were more prevalent for ultimately intentional events (26% of all questions) than ultimately unintentional ones (7%), $\chi^2(1) = 58.4$, $p < .001$, 95% CI = [.145, .240], $h = 0.53$, whereas preventability questions were more prevalent for ultimately unintentional events (18%) than ultimately intentional ones (7%), $\chi^2(1) = 22.2$, $p < .001$, 95% CI = [.062, .148], $h = 0.34$, consistent with the predicted bifurcating role of intentionality.

Conditional question sequences. The trials placed participants into information states that the Path Model predicts to be at different stages of the search sequence. This allowed a multistep tracking of the information participants sought. For example, when offered a sparse negative event description (Outcome events), participants should initially ask about agent causality, then (if causality is confirmed) about intentionality, then (if intentional) about reasons, and so on. Likewise, when offered a description of an agent-caused outcome (Causality events), they should initially ask about intentionality and so on. Below, we organize the results by the initial information state into which participants were placed.

In Outcome events (collapsing across ultimate intentionality), 61% of people's first questions targeted causality, far exceeding the next most prevalent category (intentionality: 20%), $\chi^2(1) = 42.1$, $p < .001$, 95% CI = [.295, .518], $h = 0.86$. After an initial causality question, 42% of subsequent questions targeted intentionality, exceeding the next most prevalent category (event clarification: 14%), all $\chi^2(1) > 12.4$, ps < .001, 95% CI = [.131, .423], hs > 0.64. These patterns are depicted in Figure 3, which presents the overall percentages of question types in each question position and the conditional probabilities between them (e.g., in this case, that, given an initial causality question, a subsequent intentionality question is especially likely). Among Outcome events that were ultimately intentional events, reasons questions were more prevalent in subsequent positions (19% of third and 27% of fourth questions) than were preventability questions (fewer than 10% of all questions in either position), $\chi^2(1) = 8.32$, $p < .01$. Among Outcome events that were ultimately unintentional events, preventability questions were more prevalent (12% of third and 19% of fourth questions) than were reasons questions (fewer than 5% of all questions in either position), $\chi^2(1) = 6.97$, $p < .01$.

In Causality events, 55% of first questions targeted intentionality, far exceeding the next most prevalent category (preventability: 13%), all $\chi^2(1) > 44.9$, ps < .001, 95% CI = [.303, .520], hs > 0.93. Among ultimately intentional behaviors, 57% of immediately subsequent questions targeted reasons,

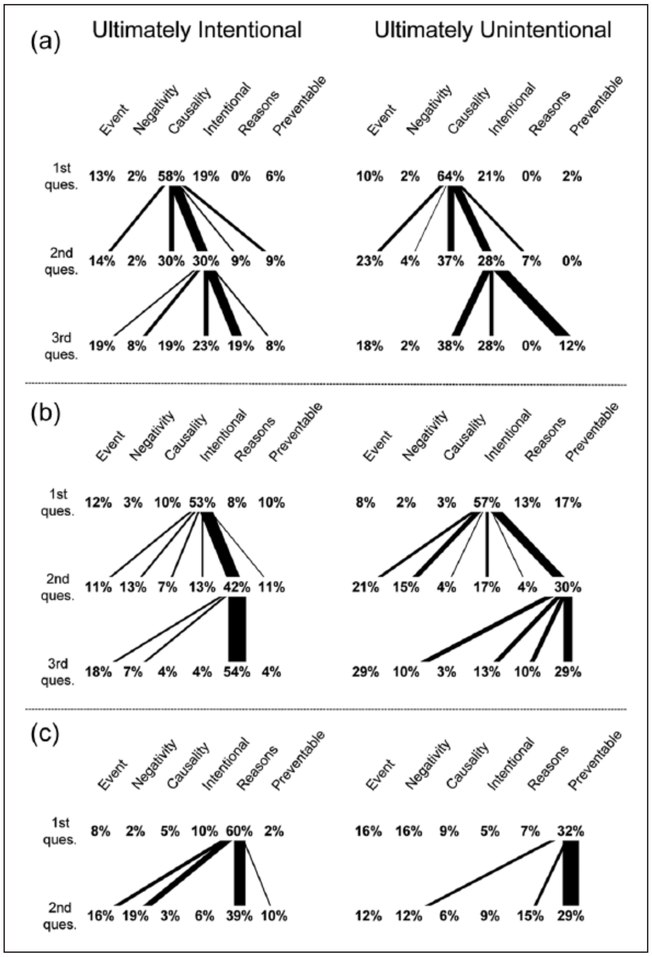

Figure 3. Proportion of questions from each category among (a) Outcome, (b) Causality, and (c) Intentionality events, for ultimately intentional (left column) and ultimately unintentional (right column) cases in Study 2. Line thickness represents conditional probability magnitude.

exceeding any other category (fewer than 10% for any other single category), all $\chi^2(1) > 11.96$, ps $< .001$, hs > 1.06. Among ultimately unintentional behaviors, 39% of immediately subsequent questions targeted preventability, significantly higher than for any other category, all $\chi^2(1) > 4.0$, ps $< .05$, hs > 0.60, except negativity (22%), $p = .19$, $h = 0.37$.

In Intentional events, 60% of first questions and 39% of second questions targeted reasons, far exceeding the next most prevalent category (event clarification) in these two positions (8% and 16%, respectively), $\chi^2(1) = 25.8$, $p < .001$. Moreover, questions about preventability were infrequent (fewer than 11% of questions in either position), as were stepping-back questions about causality and intentionality (each also fewer than 11% in either position). In Unintentional trials, 32% of first questions and 29% of second questions targeted preventability, exceeding the next most prevalent categories (event clarification [16% and 12%] and negativity [also 16% and 12%]) in these positions, $\chi^2(1) = 8.65$, $p < .01$. Moreover, questions about reasons were infrequent (fewer than 16% of questions in either position), as were stepping-back questions about causality and intentionality (each fewer than 10% in either position).

Discussion

Study 2 examined the content and patterns of information acquisition by allowing participants to ask open-ended questions en route to assigning blame. The results showed systematic patterns of information acquisition that supported the Path Model. Nearly 90% of questions targeted the information components postulated by the model: event features, agent causality, intentionality, reasons, and preventability. More importantly, specific sequences of information acquisition were consistent with predictions of the Path Model but not with nonsequential accounts. Outcome events led people to seek causality information; causality events led them to seek intentionality information; and intentional events led them to seek reasons information. Unintentional events led people to seek preventability information, although this link was weaker than the preceding model-implied links, appearing to reflect people's attention to several different concerns when considering unintentional violations—not only determining preventability but also considering the potential causal involvement of other parties and considering paths toward restitution (Darley & Pittman, 2003). Finally, the results also support the predicted bifurcating role of intentionality: reasons questions were far more prevalent than preventability questions for intentional events, but the reverse was true for unintentional events.

Studies 1 and 2 together reveal consistent patterns of information acquisition across complementary contexts. These patterns provide strong support for the Path Model of Blame and against models that deny ordered information processing en route to blame. Moreover, the ordering is not just a simple sequence but shows a hierarchical dependency of specific information searches as a function of intentionality.

A shared feature of the preceding studies is that they enabled deliberative information acquisition because of their minimal time constraints (a response window of 6,000 ms in Study 1; no time limit in Study 2). Several theorists hold that moral judgments are primarily intuitive (e.g., Haidt, 2001) or that intuitive and deliberative responses are often in conflict (Greene, 2007). It is possible, then, that Studies 1 and 2 reflect idealized patterns of information acquisition that emerge only when people have ample opportunity for deliberative processing. People might show different patterns—and ones that may challenge the Path Model—when deliberative processing is difficult or impossible (Alicke, 2014; Schein & Gray, 2014). If so, people should show less evidence of systematic hierarchical information acquisition when they must rely heavily on intuitive processing. Given an oft-cited conceptual definition of such processing as occurring "quickly, effortlessly, and automatically, such that the outcome but not the process is accessible to consciousness" (Haidt, 2001, p. 818), Study 3 implemented a tight time constraint (2 s) so as to permit quick and effortless processes but whose outcomes are still consciously reportable.

Study 3

To examine the role of deliberative processing in information acquisition toward blame, we use the information selection paradigm to contrast a minimally constrained response window (10,000 ms), permitting ample deliberative processing, with a tightly constrained window (2,000 ms), favoring fast and unreflected processing.

Method

Participants and procedure. Undergraduate participants (n = 79), who received course credit, completed the study in individual computer rooms. As in Study 1, participants completed three practice and 32 experimental trials, and then answered a brief demographic questionnaire.

Design and materials. The materials and design were identical to Study 1 except that participants were randomly assigned to either an unspeeded condition (10,000 ms to respond to each information offer; n = 42) or a speeded condition (2,000 ms; n = 37). Those in the speeded condition were told that the offers would appear only briefly, and that they would therefore need to respond as quickly as possible.

Results

Data preparation. As in Study 1, all trials (N = 3,852 trials across 79 participants) were screened for outliers. First, all responses faster than 400 ms or slower than 6,000 ms were removed (2.5% of all trials; these cutoff values afforded direct comparability to Study 1). Second, within-subject outliers were excluded (responses exceeding within-subject mean ± 2.5 within-subject SDs; 1.9% of all trials). This process yielded 3,681 valid trials for further analysis.

Acceptance response times were faster overall in the speeded condition (M = 967, SD = 346) than in the unspeeded condition (M = 1,655, SD = 925), $t(2595)$ = 25.8, p < .001, d = .95, confirming that the time pressure manipulation indeed forced participants to respond more quickly. The average acceptance response time in the unspeeded condition did not differ from that in Study 1 (M = 1,700, SD = 873), $t(3351)$ = 1.42, p > .10.

Next vs. Back. Replicating Study 1, a binomial multilevel model treating trials as the unit of analysis and including random intercepts for subjects revealed that participants in the unspeeded condition accepted Next offers more often (90%) than Back offers (25%), z = 14.52, p < .001, h = 1.45. The same pattern held in the speeded condition (89% vs. 30%), z = 12.84, p < .001, h = 1.30. As in Study 1, the Next versus Back effect was stronger for causality information (96% vs. 28%, collapsing across speed conditions; h = 1.62) than for intentionality information (83% vs. 25%; h = 1.24)—this Offer Type × Information Interaction emerged in both the unspeeded condition, z = 2.36, p < .05 and the speeded condition, z = 3.44, p < .001.

In the unspeeded condition, Next offers were accepted more quickly (M = 1,587, SD = 895) than Back offers (M = 2,144, SD = 1,136), $t(813)$ = 4.90, p < .001, 95% CI = [368.3, 861.0], d = .61, again replicating Study 1. The same pattern emerged in the speeded condition: Next acceptances were faster (M = 939, SD = 339) than Back acceptances (M = 1,017, SD = 396), $t(664)$ = 1.90, p < .06, 95% CI = [−3.0, 182.0], d = .23. The Next versus Back differentiation was, however, weaker in the speeded condition, as the Offer Type × Condition Interaction shows, $t(1477)$ = 2.12, p < .05.

Next vs. Jump. Consistent with Study 1, people in the Outcome state were more likely to accept Next offers (for causality information) than Jump offers; this was true in both the unspeeded condition (96% vs. 87%), z = 3.86, p < .001, 95% CI = [.045, .135], h = 0.33, and the speeded condition (95% vs. 81%), z = 4.95, p < .001, 95% CI = [.088, .198], h = 0.45. In contrast to Study 1, in neither condition were people in the *Cause* state any more likely to accept Next offers than Jump offers, both zs < 1.35, ps > .18, hs < 0.14.

Next offers were accepted more quickly than Jump offers in both the unspeeded (M = 1587, SD = 895 vs. M = 1,710, SD = 898), $t(1147)$ = 2.34, p < .05, 95% CI = [19.8, 224.1], d = .14, and speeded conditions (M = 939, SD = 339 vs. M = 987, SD = 345), $t(927)$ = 2.37, p < .05, 95% CI = [8.7, 92.8], d = .14. Examining the more specific comparisons provided further support for this pattern (see Figure 4). In both conditions, at the Outcome state, people accepted Next offers (for causality information) more quickly than Jump offers: unspeeded (M = 1,542, SD = 807 vs. M = 1,734, SD = 912), $t(515)$ = 2.82, p < .01, 95% CI = [59.2, 330.0], d = .22; speeded (M = 919, SD = 336 vs. M = 1,006, SD = 359), $t(423)$ = 3.44, p < .001, 95% CI = [44.1, 161.3], d = .25. Similarly, in both conditions, at the *Cause* state, people accepted Next offers (for intentionality information) more quickly than Jump offers: unspeeded (M = 1512, SD = 876 vs. M = 1664, SD = 874), $t(369)$ = 1.93, p < .06, 95% CI = [−3.1, 331.1], d = .17; speeded (M = 879, SD = 324 vs. M = 953, SD = 316), $t(290)$ = 2.04, p < .05, 95% CI = [2.4, 133.9], d = .23. The one anomalous finding is the same as in Study 1: Preventable Next offers were accepted unusually slowly as compared with Preventable Jump offers. This difference was present but nonsignificant in the unspeeded condition (M = 2,034 vs. M = 1,813), $t(144)$ = 1.58, p = .12, 95% CI = [−50.8, 462.9], d = .21; it was significant in the speeded condition (M = 1,134 vs. M = 993), $t(132)$ = 3.30, p < .01, 95% CI = [62.0, 247.4], d = .42. We explore this pattern in the General Discussion section.

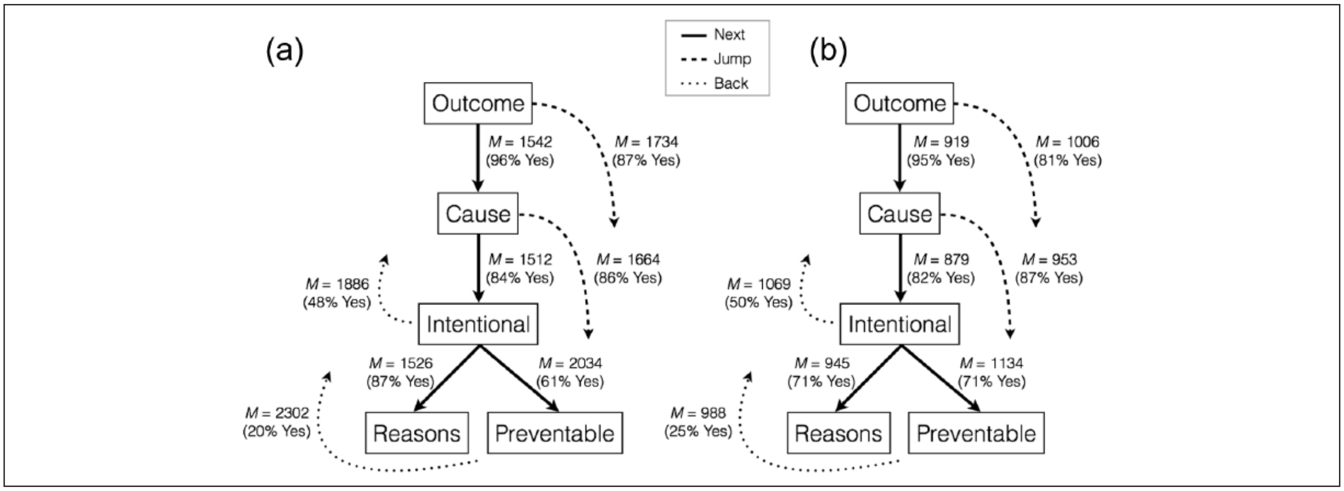

Figure 4. Average acceptance response times (and acceptance response rates) from each information state as a function of offer type in the (a) unspeeded and (b) speeded conditions (Study 3).

Discussion

Study 3 added a time pressure manipulation to Study 1's information selection paradigm to determine whether information-acquisition strategies differ under conditions of deliberative versus intuitive processing. Overall response times differed markedly between conditions, confirming that time pressure pushed people to respond much more quickly than they did in the baseline condition. Nonetheless, the results were strikingly similar across conditions (and largely consistent with those of Study 1), revealing that the patterns posited by the Path Model can emerge via either deliberative or intuitive processing. In fact, people responded quickly even in the unspeeded condition (in less than 2 s on average), suggesting that the extent of deliberation in accepting information offers is minimal.

As in Study 1, people frequently accepted Jump offers. The critical finding in support of the Path Model, however, is that people consistently took longer to accept these Jump offers than Next offers from the same information state. Thus, even though people commonly accepted offers that jumped-over components, making this mental jump took additional time.

General Discussion

Existing moral judgment models typically examine either the particular information components of moral judgments (e.g., Cushman, 2008; Schlenker et al., 1994; Shaver, 1985; Weiner, 1995) or their underlying psychological processes (e.g., Greene, 2007; Haidt, 2001). The current studies aimed to integrate these approaches by examining information acquisition, a process that organizes the critical information components en route to moral judgments. We specifically tested predictions made by the Path Model of Blame, which specifies the (a) information components that people seek out when making blame judgments, (b) the canonical sequence in which they seek these components, and (c) dependencies among the components, particularly a bifurcation of processing as a function of perceived intentionality.

Study 2 supported the first claim. When participants asked open-ended questions, the overwhelming majority of questions targeted the components identified by the Path Model (information about the norm-violating event, about causality, intentionality, reasons, and preventability), whereas questions targeting concepts such as responsibility, controllability, and character were exceedingly rare.

All three studies supported the second claim, regarding the canonical sequence of information processing. Given whatever starting information they had, people were more likely and quicker to accept offers targeting the next component in the sequence of the Path Model than information either about a prior component or one further down in the sequence (Studies 1 and 3). Even more important, when people were able to ask for whatever information they wanted (Study 2), they predominantly asked about the next component in the sequence of the Path Model. These findings are not easily accommodated by nonsequential accounts, which doubt sequential processing and therefore predict random or nonsystematic patterns of information acquisition.

Study 2 also supported the third claim, regarding the bifurcating role of intentionality. Negative events perceived to be intentional triggered a search for reasons information but rarely for preventability information. Conversely, events perceived to be unintentional triggered a search for preventability information but rarely for reasons information. However, even though people queried preventability almost exclusively for unintentional events, such events also frequently led to queries about details of the event or about other causal agents. Such queries may reflect the fact that unintentional negative behaviors are less scripted and less expected than intentional ones (Rosset, 2008), as the absence of equifinality (Heider, 1958) widens the search for potential causes. The processing of unintentional event information is thus more complex than a straightforward unintentional → preventability link would suggest.

The information selection paradigm revealed an additional surprising pattern regarding preventability information. Studies 1 and 3 both showed that people were quick to accept preventability information before yet learning about intentionality (Jump offers) but were slow to do so when they knew that the event was unintentional (Next offers). This pattern might suggest that people's lay understanding of the term preventability might be akin to controllability, or the extent to which the agent could control or foresee the outcome (Alicke, 2000; Weiner, 1995). As controllability can apply to both intentional and unintentional behavior, people might have accepted preventability information as a way of learning about intentionality. Importantly though, the more open-ended design of Study 2 revealed that people typically sought intentionality information before preventability information, and they sought specific preventability information (e.g., about carelessness or recklessness) once they had learned that an event was unintentional.

Last, Study 3 showed that information-acquisition strategies were highly consistent under both low time constraint (allowing ample deliberative processing) and high time constraint (demanding intuitive processing). In both cases, given their current knowledge state, people were most inclined to seek information targeting the next component as predicted by the Path Model. These findings further demonstrate the robust and systematic nature of information acquisition in blame judgments, and they suggest that the pathways specified by the Path Model can be actualized across different modes of processing.

Implications for Models of Moral Judgment

The results revealed the importance of various concepts that have been highlighted, in different numbers and combinations, by previous models, including causality (Alicke, 2000), intentionality (Shaver, 1985), preventability (Schlenker et al., 1994; Weiner, 1995), and reasons (Cushman, 2008). In this sense, the current findings are consistent with some general claims of these models. Only the Path Model, however, integrates all the relevant components into one framework. Moreover, previous models are typically agnostic or doubtful about hierarchies between the relevant components and therefore do not predict differential acceptance speeds as revealed in Studies 1 and 3, nor specific contingencies in spontaneous question-asking as revealed in Study 2. Furthermore, none of the previous models predict the bifurcating influence of intentionality on information acquisition, as supported by the current studies.

Much prior theorizing argues that all, or at least most, moral judgments tend to be unreflective, fast, and intuitive (Greene, 2007; Haidt, 2001), and some scholars have suggested that a hierarchical conceptual structure such as the one posited by the Path Model simply could not guide such intuitive processing (Alicke, 2014; Schein & Gray, 2014). But people's information-acquisition behavior was remarkably consistent in Study 3 regardless of time pressure, indicating that processing en route to blame is highly systematic and organized, integrating many pieces of information, though nonetheless fast and (perhaps) intuitive. Such systematic and complex information integration has often been ascribed only to slow and deliberative processes. We believe, however, that the contrast between intuitive and deliberative processes is less fruitful than the degree of concept- and information richness (cf. Keren & Schul, 2009; Kruglanski & Gigerenzer, 2011). Blame judgments require the moral perceiver to process and integrate considerable information (Malle et al., 2014), whereas wrongness or badness judgments require less integration and have weaker social regulatory demands, thereby likely making them faster and more unreflective than blame (Malle et al., 2014). And all moral judgments—blame included—are no doubt faster and less plodding than, say, mental arithmetic, which is perhaps a prototype of deliberative reasoning.

Future Directions

Motivated information search. The information-acquisition paradigms implemented here offer useful methods for examining other topics in the literature on moral judgment. Several scholars have suggested that perceivers often engage in a motivated or biased search for blame-consistent information, so as to cast the target as particularly blameworthy (Alicke, 2000; Ames & Fiske, 2013). Accordingly, rather than establishing the truth values of key information components, perceivers should be inclined to confirm that these components have blame-implying values. For example, upon discovering a particularly negative event or agent, perceivers should prefer intentionality-implying information over unintentionality-implying information. They should likewise prefer information that suggests exacerbating rather than mitigating reasons. These predictions would be consistent with general findings of confirmation bias (Kunda, 1990; Nickerson, 1998), as well as with specific findings from legal decision-making contexts, in which people seek and interpret evidence in a guilt-presumptive fashion (Ask & Granhag, 2005; Kassin, Goldstein, & Savitsky, 2003).[6] The present paradigm could test such predictions of motivated information search perhaps more directly than previous paradigms, because it allows direct measurement of participants' information-seeking behavior.

Emotional influences on information search. Relatedly, a substantial body of evidence suggests that emotions play a pivotal role in shaping moral judgments (Greene, 2007; Haidt, 2001; Horberg et al., 2009; Moretto et al., 2010), and future research can specifically explore the role of emotion in information-acquisition behavior. As Pizarro, Inbar, and Helion (2011) note, there are many potential explanations for the operative mechanism linking emotion to moral judgment. Emotion might be a "moralizer," turning nonmoral violations into moral ones, or an "amplifier," making already moral violations seem more extreme (Pizarro et al., 2011). The latter process may have several manifestations. As a "judgment amplifier," emotion would directly influence moral judgments without altering the mediating information processing (about causality, intentionality, etc.). As a "component amplifier," emotion would influence perceptions of specific information components, which thereby influence moral judgments (cf. Alicke, 2000; Feigenson & Park, 2006). Anger may exhibit both patterns: amplifying the desire for punishment (Goldberg, Lerner, & Tetlock, 1999) and influencing punishment via its effect on intention inferences (Ask & Pina, 2011).

However, standard vignette studies have not yet fully addressed the timing and sequence of affective and moral information processing. Information-acquisition paradigms could evaluate the relative impact of these mechanisms by assessing how various negative emotions may lead perceivers to ask particular questions or accept particular information en route to making moral judgments. If negative emotions are "judgment amplifiers," they should have minimal impact on information search but nonetheless elicit more severe blame. If they are "component amplifiers," negative emotions should influence the information search process itself.

Information acquisition in naturalistic contexts. The present studies present laboratory-based evidence for systematic patterns of information acquisition. Will such patterns likewise emerge in more naturalistic contexts? Such contexts undoubtedly involve a great deal of additional complexity, such as multiple potential sources of information, partial or competing information, or information presented simultaneously. Future research will have to examine these aspects, but there nonetheless are hints that information acquisition in naturalistic contexts proceeds in a similar fashion to the patterns revealed here.

In December 2016, a fire ripped through an Oakland warehouse called the Ghost Ship, killing 36 people. *The New York (NY) Times* ("The Oakland Fire," 2016), in its coverage, presented a snapshot of its information-acquisition process, first identifying the negative event at hand ("A Devastating Fire. Thirty-six Lives Lost") and then posing various questions that its journalists would be investigating. Notably, many of these questions directly map onto the major elements of the Path Model, including causality ("What was the cause of the blaze?"), preventability ("Could the city have done more to prevent the Ghost Ship fire?"), and of course blame or responsibility itself ("Who is ultimately responsible?"). Although far from comprehensive, this example of naturalistic information acquisition is consistent with the general concepts and claims of the Path Model.

However, naturalistic search strategies may often proceed in more nuanced ways, rather than strictly or solely along the steps of the Path Model. First, perceivers might receive or acquire information about other blame-relevant features (e.g., about an agent's character or past behavior) or seek to clarify components of the Path Model by obtaining more fine-grained information (e.g., about an agent's effort or planning as clues about the intentionality of the behavior). Second, some negative events afford strong inferences about certain information components (e.g., intentionality) and thereby obviate the need to search explicitly for such information. Assault, for example, is almost certainly intentional, whereas fire blazes are often caused unintentionally. Indeed, in its investigation of the Oakland fire, *The NY Times* did not include an explicit question about the intentionality of the event, instead assuming it was unintentional. Third, when confronted with actual negative events, people aim not just to assess blame but also to understand relevant historical or contextual influences and to consider strategies for preventing similar events in the future. Many of *The NY Times*' questions went beyond the scope of blaming agents (and thus beyond the Path Model), presenting information about the victims and exploring implications for future policy and regulation.

Perceivers' own ideological or personal commitments may also guide their information acquisition, leading them to seek or interpret certain information in preferred ways. People who endorse binding moral foundations—including loyalty, purity, and obedience to authority—are more likely to see victims as responsible and blameworthy for their own suffering (Niemi & Young, 2016). Such individuals might therefore actively seek information that causally implicates victims or might interpret a given piece of information as providing greater evidence for the victim's causal role. People who oppose abortion are more likely to view an instance of abortion as an active *killing* rather than as a passive *allowing* to die (Cushman, Knobe, & Sinnott-Armstrong, 2008); consequently, when considering such cases, they might already assume information about intentionality and unjustified motives. In these cases, the ordered information search may be intact, but the norms that are seen as violated by the original event, and the candidate causes assumed for the event, differ as a function of ideology or prior convictions (see Malle et al., 2014, for further discussion of such "presets" of information).

Computational processes. The present studies reveal that people engage in systematic patterns of information acquisition when assessing blame, consistent with the claims of the Path Model. As with any claim about psychological processing, it will be important for future research to better understand such patterns at a computational level. What are the specific computational and neurological mechanisms that track, represent, and integrate the relevant components of moral judgment? Initial evidence for the neural underpinnings of moral judgment has been available for a while (e.g., Greene, Sommerville, Nystrom, Darley, & Cohen, 2001; Van Berkum, Holleman, Nieuwland, Otten, & Murre, 2009), but scholars have only recently begun working to specify corresponding formal computational models (e.g., Crockett, 2016). Computational modeling of moral judgment will be a fruitful direction for future research, but so too will a computational analysis of domain-general social cognitive processes. Moral judgment in general, and the Path Model specifically, draws on fundamental processes of theory of mind, attribution, and affective processing (Guglielmo, Monroe, & Malle, 2009; Young & Dungan, 2012). Heightened understanding of these basic social cognitive building blocks will likewise enhance our understanding of moral cognition and its underlying processes.

Summary

The current studies examined information acquisition in blame judgment via the use of two novel paradigms: one that offered participants particular types of information and one that allowed them to seek whichever information they wished to acquire. The Path Model of Blame posits that perceivers process information about negative events in a canonical order, first clarifying information about causality, then intentionality, then either reasons (for intentional events) or preventability (for unintentional ones). These systematic patterns of information acquisition emerged across both methodological paradigms, and even when people faced severe time pressure. The information-acquisition paradigms presented here offer promising directions for future research, particularly in extending research on motivated or emotionally guided processing. The literature will benefit from continued attention to information acquisition in moral judgment; there is a great deal more searching to be done.

Acknowledgments

Thanks to Steven Sloman, Joachim Krueger, Andrew Monroe, and Joanna Korman for their helpful discussion and feedback, and to Nicole Boucher for her diligent coding work in Study 2.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was partially supported by an APA Dissertation Research Award.

Notes

1. The Path Model also notes the importance of *obligation*; however, because this component is rarely examined in the literature and can typically be assumed to be present, it will not be investigated here.
2. In trials with multiple offers, offer types were defined based on the participant's information status after the previous offer. For example, for an *Outcome* event, if the participant accepted an initial CAUSE cue, a subsequent INTENTIONAL cue would constitute a Next offer.
3. If it was not a yes/no question or if the database did not contain the answer, the experimenter sent a scripted prompt that told the participant to ask a different question (this occurred for only 4.4% of questions).
4. The categories of *response* (4.2% of all questions), *uncodable* (1.6%) and *obligation* (.6%), were infrequent and were therefore omitted from Table 1 and Figure 3.
5. Participants might have refrained from asking about responsibility because they felt that their task of assigning blame was identical to attributing responsibility. However, this assumption would be inconsistent with evidence that these two judgments are distinct (e.g., Harvey & Rule, 1978). Moreover, when participants used the word "responsibility" in their questions, they referred to various meanings, including *causal responsibility* (e.g., "Was the cause of her not paying attention her own responsibility?"), *general moral responsibility* (e.g., "Is one person responsible for this?"), *obligation* (e.g., "Was there one individual whose responsibility it was . . . ?"), and *admission* (e.g., "Did she take responsibility for her actions?").
6. There is much theoretical debate over whether such patterns of hypothesis confirmation should properly be considered *biases*. Some scholars have argued, rather, that they often reflect normative, rational Bayesian inferences (Fischhoff & Beyth-Marom, 1983).

Supplemental Material

The online supplemental material is available with the manuscript on the PSPB website.

References

Alicke, M. D. (2000). Culpable control and the psychology of blame. *Psychological Bulletin*, *126*, 556-574.

Alicke, M. D. (2014). Evaluating blame hypotheses. *Psychological Inquiry*, *25*, 187-192.

Alicke, M. D., Weigold, M. F., & Rogers, S. L. (1990). Inferring intentions and responsibility from motives and outcomes: Evidential and extra-evidential judgments. *Social Cognition*, *8*, 286-305.

Ames, D. L., & Fiske, S. T. (2013). Intentional harms are worse, even when they're not. *Psychological Science*, *24*, 1755-1762.

Ask, K., & Granhag, P. A. (2005). Motivational sources of confirmation bias in criminal investigations: The need for cognitive closure. *Journal of Investigative Psychology and Offender Profiling*, *2*, 43-63.

Ask, K., & Pina, A. (2011). On being angry and punitive: How anger alters perception of criminal intent. *Social Psychological and Personality Science*, *2*, 494-499.

Crockett, M. J. (2016). How formal models can illuminate mechanisms of moral judgment and decision making. *Current Directions in Psychological Science*, *25*, 85-90.

Cushman, F. (2008). Crime and punishment: Distinguishing the roles of causal and intentional analyses in moral judgment. *Cognition*, *108*, 353-380.

Cushman, F., Knobe, J., & Sinnott-Armstrong, W. (2008). Moral appraisals affect doing/allowing judgments. *Cognition*, *108*, 281-289.

Darley, J. M., & Pittman, T. S. (2003). The psychology of compensatory and retributive justice. *Personality and Social Psychology Review*, *7*, 324-336.

Dhami, M. K., & Harries, C. (2010). Information search in heuristic decision making. *Applied Cognitive Psychology*, *24*, 571-586.

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*, 175-191.

Feigenson, N., & Park, J. (2006). Emotions and attributions of legal responsibility and blame: A research review. *Law and Human Behavior*, *30*, 143-161.

Fischhoff, B., & Beyth-Marom, R. (1983). Hypothesis evaluation from a Bayesian perspective. *Psychological Review*, *90*, 239-260.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*, 451-482.

Goldberg, J. H., Lerner, J. S., & Tetlock, P. E. (1999). Rage and reason: The psychology of the intuitive prosecutor. *European Journal of Social Psychology*, *29*, 781-795.

Gray, K., Young, L., & Waytz, A. (2012). Mind perception is the essence of morality. *Psychological Inquiry*, *23*, 101-124.

Greene, J. D. (2007). The secret joke of Kant's soul. In W. Sinnott-Armstrong (Ed.), *Moral Psychology: Vol. 3. The neuroscience of morality: Emotion, disease, and development* (pp. 35-79). Cambridge, MA: MIT Press.

Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, *293*, 2105-2108.

Guglielmo, S. (2015). Moral judgment as information processing: An integrative review. *Frontiers in Psychology*, *6*, Article 1637.

Guglielmo, S., Monroe, A. E., & Malle, B. F. (2009). At the heart of morality lies folk psychology. *Inquiry*, *52*, 449-466.

Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, *108*, 814-834.

Harvey, M. D., & Rule, B. G. (1978). Moral evaluations and judgments of responsibility. *Personality and Social Psychology Bulletin*, *4*, 583-588.

Heider, F. (1958). *The psychology of interpersonal relations*. New York, NY: John Wiley.

Horberg, E. J., Oveis, C., Keltner, D., & Cohen, A. B. (2009). Disgust and the moralization of purity. *Journal of Personality and Social Psychology*, *97*, 963-976.

Inbar, Y., Pizarro, D. A., & Cushman, F. (2012). Benefitting from misfortune: When harmless actions are judged to be morally blameworthy. *Personality and Social Psychology Bulletin*, *38*, 52-62.

Jacoby, J., Jaccard, J., Kuss, A., Troutman, T., & Mazursky, D. (1987). New directions in behavioral process research: Implications for social psychology. *Journal of Experimental Social Psychology*, *23*, 146-175.

Kassin, S. M., Goldstein, C. C., & Savitsky, K. (2003). Behavioral confirmation in the interrogation room: On the dangers of presuming guilt. *Law and Human Behavior*, *27*, 187-203.

Keren, G., & Schul, Y. (2009). Two is not always better than one: A critical evaluation of two-system theories. *Perspectives on Psychological Science*, *4*, 533-550.

Kruglanski, A. W., & Gigerenzer, G. (2011). Intuitive and deliberative judgments are based on common principles. *Psychological Review*, *118*, 97-109.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*, 480-498.

Lagnado, D. A., & Channon, S. (2008). Judgments of cause and blame: The effects of intentionality and foreseeability. *Cognition*, *108*, 754-770.

Malle, B. F., Guglielmo, S., & Monroe, A. E. (2012). Moral, cognitive, and social: The nature of blame. In J. Forgas, K. Fiedler, & C. Sedikides (Eds.), *Social thinking and interpersonal behavior* (pp. 311-329). Philadelphia, PA: Psychology Press.

Malle, B. F., Guglielmo, S., & Monroe, A. E. (2014). A theory of blame. *Psychological Inquiry*, *25*, 147-186.

Moretto, G., Làdavas, E., Mattioli, F., & di Pellegrino, G. (2010). A psychophysiological investigation of moral judgment after ventromedial prefrontal damage. *Journal of Cognitive Neuroscience*, *22*, 1888-1899.

Nelson, S. A. (1980). Factors influencing young children's use of motives as moral criteria. *Child Development*, *51*, 823-829.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, *2*, 175-220.

Niemi, L., & Young, L. (2016). When and why we see victims as responsible: The impact of ideology on attitudes toward victims. *Personality and Social Psychology Bulletin*, *42*, 1227-1242.

The Oakland fire: Delving into what happened, and why. (2016, December 8). *The New York Times*. Retrieved from https://www.nytimes.com/2016/12/08/us/the-oakland-fire-delving-into-what-happened-and-why.html

Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational Behavior and Human Performance*, *16*, 366-387.

Pizarro, D. A., Inbar, Y., & Helion, C. (2011). On disgust and moral judgment. *Emotion Review*, *3*, 267-268.

Pizarro, D. A., & Tannenbaum, D. (2011). Bringing character back: How the motivation to evaluate character influences judgments of moral blame. In M. Mikulincer & P. R. Shaver (Eds.), *The social psychology of morality: Exploring the causes of good and evil* (pp. 91-108). Washington, DC: American Psychological Association Press.

Reisen, N., Hoffrage, U., & Mast, F. W. (2008). Identifying decision strategies in a consumer choice situation. *Judgment and Decision Making*, *3*, 641-658.

Rosset, E. (2008). It's no accident: Our bias for intentional explanations. *Cognition*, *108*, 771-780.

Schein, C., & Gray, K. (2014). The prototype model of blame: Freeing moral cognition from linearity and little boxes. *Psychological Inquiry*, *25*, 236-240.

Schlenker, B. R., Britt, T. W., Pennington, J., Murphy, R., & Doherty, K. (1994). The triangle model of responsibility. *Psychological Review*, *101*, 632-652.

Shaver, K. G. (1985). *The attribution of blame: Causality, responsibility, and blameworthiness*. New York, NY: Springer.

Sloman, S. A., Fernbach, P. M., & Ewing, S. (2009). Causal models: The representational infrastructure for moral judgment. In B. H. Ross (Series Ed.),D. M. Bartels, C. W. Bauman, L. J. Skitka, & D. L. Medin (Eds.), *Psychology of Learning and Motivation: Vol. 50. Moral judgment and decision making* (pp. 1-26). San Diego, CA: Academic Press.

Van Berkum, J. J. A., Holleman, B., Nieuwland, M., Otten, M., & Murre, J. (2009). Right or wrong? The brain's fast response to morally objectionable statements. *Psychological Science*, *20*, 1092-1099.

Weiner, B. (1995). *Judgments of responsibility: A foundation for a theory of social conduct*. New York, NY: Guilford Press.

Woolfolk, R. L., Doris, J. M., & Darley, J. M. (2006). Identification, situational constraint, and social cognition: Studies in the attribution of moral responsibility. *Cognition*, *2*, 283-301.

Young, L., & Dungan, J. (2012). Where in the brain is morality? Everywhere and maybe nowhere. *Social Neuroscience*, *7*, 1-10.

Young, L., & Saxe, R. (2009). Innocent intentions: A correlation between forgiveness for accidental harm and neural activity. *Neuropsychologia*, *47*, 2065-2072.

Zelazo, P. D., Helwig, C. C., & Lau, A. (1996). Intention, act, and outcome in behavioral prediction and moral judgment. *Child Development*, *67*, 2478-2492.