Journal of Economic Perspectives—Volume 30, Number 3—Summer 2016—Pages 133–140

The Mechanics of Motivated Reasoning

Nicholas Epley and Thomas Gilovich

Whenever we see voters explain away their preferred candidate's weaknesses, dieters assert that a couple scoops of ice cream won't *really* hurt their weight loss goals, or parents maintain that their children are unusually gifted, we are reminded that people's preferences can affect their beliefs. This idea is captured in the common saying, "People believe what they want to believe."

But people don't *simply* believe what they want to believe. The psychological mechanisms that produce motivated beliefs are much more complicated than that. Personally, we'd like to believe that our contributions to the psychological literature might someday rival those of Daniel Kahneman, but, try as we might, the disparity in citations, prizes, invitations—you name it—makes holding such a belief impossible. People generally *reason* their way to conclusions they favor, with their preferences influencing the way evidence is gathered, arguments are processed, and memories of past experience are recalled. Each of these processes can be affected in subtle ways by people's motivations, leading to biased beliefs that feel objective (Gilovich and Ross 2015; Pronin, Gilovich, and Ross 2004). As Kunda (1990) put it, "people motivated to arrive at a particular conclusion attempt to be rational and to construct a justification of their desired conclusion that would persuade a dispassionate observer. They draw the desired conclusion only if they can muster up the evidence necessary to support it" (p. 482–83). Motivated reasoning is constrained.

■ *Nicholas Epley is the John T. Keller Professor of Behavioral Science, University of Chicago, Booth School of Business, Chicago, Illinois. Thomas Gilovich is the Irene Blecker Rosenfeld Professor of Psychology, Cornell University, Ithaca, New York. Their email addresses are epley@chicagobooth.edu and tdg1@cornell.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
http://dx.doi.org/10.1257/jep.30.3.133

doi=10.1257/jep.30.3.133

Psychological research makes it clear, in other words, that "motivated beliefs" are guided by motivated reasoning—reasoning in the service of some self-interest, to be sure, but reasoning nonetheless. We hope that being explicit about what psychologists have learned about motivated reasoning will help clarify the types of motivated beliefs that people are most likely to hold, specify when such beliefs are likely to be strong and when they are likely to be relatively weak or fragile, and illuminate when they are likely to guide people's behavior.

In this introduction, we set the stage for the discussion of motivated beliefs in the papers that follow by providing more detail about the underlying psychological processes that guide motivated reasoning, including a discussion of the varied motives that drive motivated reasoning and a description of how goals can direct motivated reasoning to produce systematically biased beliefs. The first paper in this symposium, by Roland Bénabou and Jean Tirole, presents a theoretical framework for how motives might influence behavior in several important domains; two additional papers focus on specific motives that can guide motivated reasoning: Russell Golman, George Loewenstein, Karl Ove Moene, and Luca Zarri discuss how a "preference for belief consonance" leads people to try to reduce the gap between their beliefs and those of relevant others, and Francesca Gino, Michael Norton, and Roberto Weber consider how people engage in motivated reasoning to feel as if they are acting morally, even while acting egoistically.

A more detailed understanding of motivated beliefs and motivated reasoning yields a middle-ground view of the quality of human judgment and decision-making. It is now abundantly clear that people are not as smart and sophisticated as rational agent models assert (Kahneman and Tversky 2000; Thaler 1991; Simon 1956), in the sense that people do not process information in unbiased ways. But people are also not as simple-minded, naïve, and prone to simply ignoring unpalatable information as a shallow understanding (or reporting) of motivated beliefs might suggest.

Motives for Reasoning

People reason to prepare for action, and so reasoning is motivated by the goals people are trying to achieve. A coach trying to win a game thinks about an opponent's likely moves more intensely than a cheerleader trying to energize the crowd. A lawyer trying to defend a client looks for evidence of innocence, whereas a lawyer seeking to convict tries to construct a chain of reasoning that will lead to a guilty verdict. A person feeling guilty about harming another focuses on ways to assuage the guilt, while the person harmed is likely to focus on the nature and extent of the harm. As the great psychologist and philosopher William James (1890, p. 333) wrote more than a century ago: "My thinking, is first and last and always for the sake of my doing, and I can only do one thing at a time."

One of the complexities in understanding motivated reasoning is that people have many goals, ranging from the fundamental imperatives of survival and

reproduction to the more proximate goals that help us survive and reproduce, such as achieving social status, maintaining cooperative social relationships, holding accurate beliefs and expectations, and having consistent beliefs that enable effective action. Sometimes reasoning directed at one goal undermines another. A person trying to persuade others about a particular point is likely to focus on reasons why his arguments are valid and decisive—an attentional focus that could make the person more compelling in the eyes of others but also undermine the accuracy of his assessments (Anderson, Brion, Moore, and Kennedy 2012). A person who recognizes that a set of beliefs is strongly held by a group of peers is likely to seek out and welcome information supporting those beliefs, while maintaining a much higher level of skepticism about contradictory information (as Golman, Loewenstein, Moene, and Zarri discuss in this symposium). A company manager narrowly focused on the bottom line may find ways to rationalize or disregard the ethical implications of actions that advance short-term profitability (as Gino, Norton, and Weber discuss in this symposium).

The crucial point is that the process of gathering and processing information can systematically depart from accepted rational standards because one goal—desire to persuade, agreement with a peer group, self-image, self-preservation—can commandeer attention and guide reasoning at the expense of accuracy. Economists are well aware of crowding-out effects in markets. For psychologists, motivated reasoning represents an example of crowding-out in attention.

In any given instance, it can be a challenge to figure out which goals are guiding reasoning. Consider the often-cited examples of "above-average" effects in self-evaluation: on almost any desirable human trait, from kindness to trustworthiness to the ability to get along with others, the average person consistently rates him- or herself above average (Alicke and Govorun 2005; Dunning, Meyerowitz, and Holzberg 1989; Klar and Giladi 1997). An obvious explanation for this result is that people's reasoning is guided by egoism, or the goal to think well of oneself. Indeed, a certain percentage of above-average effects can be explained by egoism because unrelated threats to people's self-image tend to increase the tendency for people to think they are better than others, in an apparent effort to bolster their self-image (as in Beauregard and Dunning 1998).

But above-average effects also reflect people's sincere attempts to assess accurately their standing in the world. For instance, many traits are ambiguous and hard to define, such as leadership or creativity. When people try to understand where they stand relative to their peers on a given trait, people quite naturally focus on what they know best about that trait—and what they know best are the personal strengths that guide their own lives. As Thomas Schelling (1978, pp. 64–65) put it, "Careful drivers give weight to care, skillful drivers give weight to skill, and those who think that, whatever else they are not, at least they are polite, give weight to courtesy, and come out high on their own scale. This is the way that every child has the best dog on the block." The above-average effect, in other words, can result from a self-enhancement goal, or from a non-motivated tendency to define traits egocentrically. Supporting Schelling's analysis, the above-average effect is significantly

reduced when traits are given precise definitions, or when the traits are inherently less ambiguous such as "punctual" or "tall" (Dunning, Meyerowitz, and Holzberg 1989).

Knowing which goal is guiding reasoning is critical for predicting the influence of specific interventions. For example, economists routinely predict that biases in judgment will be reduced when the stakes for accurate responding are high. This prediction implicitly assumes that people are not trying to be accurate already. But in fact, many cognitive biases are not affected by increased incentives for accuracy because the individuals in question are already trying hard to be accurate (Camerer and Hogarth 1999). Increasing the incentive to achieve a goal should influence behavior only when people are not already trying to achieve that goal.

How Motives Influence Beliefs

Understanding that multiple goals can shape reasoning does not explain *how* reasoning can become systematically biased. Reasoning involves the recruitment and evaluation of evidence. Goals can distort both of these basic cognitive processes.

Recruiting Evidence

When recruiting evidence to evaluate the validity of a given belief, an impartial judge would consider all of the available evidence. Most people do not reason like impartial judges, but instead recruit evidence like attorneys, looking for evidence that supports a desired belief while trying to steer clear of evidence that refutes it. In one memorable example, essayist Johanna Gohmann (2015) describes her improbable teenage crush on the actor Jimmy Stewart, and her reaction as she learned more and more about Mr. Stewart: "As I flipped through the pages my eyes skimmed words like 'womanizer' and 'FBI informant,' and I slapped it shut, reading no further." If you avoid recruiting evidence that you would prefer not to believe, your beliefs will be based on only a comforting slice of the available facts. One prominent example of motivated avoidance comes from studies of people's reactions to the prospect of having Huntington's disease: few people who are at risk of getting the disease get tested before showing symptoms, and those with symptoms who avoid testing have beliefs that are just as optimistic as those who show no symptoms (Oster, Shoulson, and Dorsey 2013).

Even when people do not actively avoid information, psychological research consistently demonstrates that they have an easier time recruiting evidence supporting what they *want* to be true than evidence supporting what they want to be false. But even here, people are still responsive to reality and don't simply believe whatever they want to believe. Instead, they recruit subsets of the relevant evidence that are biased in favor of what they want to believe. Failing to recognize the biased nature of their information search leaves people feeling that their belief is firmly supported by the relevant evidence.

Biased information processing can be understood as a general tendency for people to ask themselves very different questions when evaluating propositions they favor versus oppose (Gilovich 1991). When considering propositions they would prefer to be true, people tend to ask themselves something like "Can I believe this?" This evidentiary standard is rather easy to meet; after all, *some* evidence can usually be found even for highly dubious propositions. Some patients will get better after undergoing even a worthless treatment; someone is bound to conform to even the most baseless stereotype; some fact can be found to support even the wackiest conspiracy theory.

In contrast, when considering propositions they would prefer *not* be true, people tend to ask themselves something like "*Must* I believe this?" This evidentiary standard is harder to meet; after all, *some* contradictory evidence can be found for almost any proposition. Not all patients benefit from demonstrably effective treatments; not all group members conform to the stereotypes of their group; even the most comprehensive web of evidence will have a few holes. More compelling evidence is therefore required to pass this "Must I?" standard. In this way, people can again end up believing what they want to believe, not through mindless wishful thinking but rather through genuine reasoning processes that seem sound to the person doing it.

In one study that supports this Can I?/Must I? distinction, students were told that they would be tested for an enzyme deficiency that would lead to pancreatic disorders later in life, even among those (like presumably all of them) who were not currently experiencing any symptoms (Ditto and Lopez 1992). The test consisted of depositing a small amount of saliva in a cup and then putting a piece of litmus paper into the saliva. Half the participants were told they would know they had the enzyme deficiency if the paper changed color; the other half were told they would know they had it if the paper *did not* change color. The paper was such that it did not change color for anyone.

Participants in these two conditions reacted very differently to the same result—the unchanged litmus paper. Those who thought it reflected good news were quick to accept that verdict and did not keep the paper in the cup very long. Those who thought the unchanged color reflected bad news, in contrast, tried to recruit more evidence. They kept the paper in the cup significantly longer, even trying out (as the investigators put it) "a variety of different testing behaviors, such as placing the test strip directly on their tongue, multiple redipping of the original test strip (up to 12 times), as well as shaking, wiping, blowing on, and in general quite carefully scrutinizing the recalcitrant . . . test strip." A signal that participants wanted to receive was quickly accepted; a signal they did not want to receive was subjected to more extensive testing.

People's motivations thus do not directly influence what they believe. Instead, their motivations guide what information they consider, resulting in favorable conclusions that seem mandated by the available evidence.

Evaluating Evidence

Of course, even when looking at the very same evidence, people with different goals can interpret it differently and come to different conclusions. In one telling

experiment cited in this symposium, participants who were randomly assigned to play the role of a prosecuting attorney judged the evidence presented in trial to be more consistent with the defendant's guilt than did participants randomly assigned to play the role of the defense attorney (Babcock and Loewenstein 1997).

These distorting influences can take many forms, influencing the apparent meaning of the evidence before us. For instance, any given action can be thought of in multiple ways. A father lifting a child off the floor could be described as "picking up a child" or "caring for the child." The two equally apt descriptions have very different meanings. Caring for a child is a more significant, benevolent act than simply picking up the child. A person trying to extol a parent's character will be more likely to code the event in a higher-level term like "caring" than a person trying to demean a parent's character. Differences in how people construe the very same action can lead two people to observe the same event but "see" very different things (Maas, Salvi, Arcuri, and Semin 1989; Trope and Lieberman 2003; Vallacher and Wegner 1987).

Psychologists have examined a host of ways in which people's goals influence how they evaluate information, and we won't review that voluminous literature here. But it is worth noting that psychologists have been especially interested in the distortions that arise in the service of consistency. Leon Festinger's (1957) theory of cognitive dissonance has been particularly influential. The central idea is that people are motivated to reconcile any inconsistencies between their actions, attitudes, beliefs, or values. When two beliefs are in conflict, or when an action contradicts a personal value, the individual experiences an unpleasant state of arousal that leads to psychological efforts to dampen or erase the discrepancy, often by changing a belief or attitude.

Festinger's (1957) theory stemmed in part from his earlier work on group dynamics and what he called "pressures to uniformity" (Festinger 1950). When differences of opinion arise within a group, a palpable tension arises that group members try to resolve. That tension, he maintained, is diminished only when agreement is achieved, typically by the majority pressuring the minority to go along. Festinger's theory of cognitive dissonance essentially took what he had observed in groups and put it in the head of the individual: that is, what plays out interpersonally in group dynamics also takes place in individual psychodynamics. We all feel psychological discomfort when our actions, attitudes, beliefs, or values conflict, and that discomfort leads us to seek ways to reduce the dissonance.

By focusing on cognitive processes that occur in the head of the individual, Festinger (1957) helped to usher in a period in which social psychology became a lot less social. But dissonance reduction is often a group effort. We help one another feel better about potentially upsetting inconsistencies in our thoughts and deeds. Our friends reassure us that we chose the right job, the right house, or the right spouse. We console an acquaintance who's messed up by saying that "it's not so bad," "he had it coming," or "things would have turned out the same regardless of what you did." Indeed, whole societies help their members justify the ill-treatment of minorities, the skewed division of resources, or the degradation of

the environment through a variety of mechanisms, including everyday discourse, mass media messages, the criminal code, and even how the physical environment is structured.

The social element of rationalization and dissonance reduction fits nicely with the insightful piece by Golman, Loewenstein, Moene, and Zarri on people's preference for belief consonance. Furthermore, by connecting the preference for belief consonance to the existing literature on dissonance reduction, a great body of empirical research can be tapped to advance our understanding of when and why people will have an easy time achieving the belief consonance they seek, and when and why they are likely to struggle.

Coda

The most memorable line from the classic film *Gone with the Wind*—indeed, the most memorable line in the history of American movies according to the American Film Institute—is Rhett Butler's dismissive comment, "Frankly Scarlett, I don't give a damn." But a different line from that film has attracted more interest from psychologists: Scarlett O'Hara's frequent lament, "I can't think about that right now. . . . I'll think about it tomorrow."

The comment captures people's intuitive understanding of how motivations and emotions influence our judgments and decisions. When Scarlett doesn't want to accept some unwelcome possibility, she willfully cuts herself off from the relevant evidence. She can continue to believe what she wants because she never consults evidence that would lead her to believe differently.

Scarlet's path is one way that people can end up believing what they want to believe. But as we have noted, there are many others. Furthermore, people's preferred beliefs, developed and sustained through whatever path, guide their behavior whenever they are called to mind as choices are made. The path from motives to beliefs to choices should not be a black box to be filled with analytically convenient assumptions. Different motives can guide reasoning in different ways on different occasions—altering how information is recruited and evaluated—depending on what a person is preparing to do. We are delighted to see a topic with such a long history in psychological science being taken seriously by economists.

■ *Thanks to George Loewenstein, who took the leading role in stimulating and organizing the papers that appear in this symposium.*

References

Alicke, Mark D., and Olesya Govorun. 2005. "The Better-than-Average Effect." In *The Self in Social Judgment*, edited by Mark D. Alicke, David A. Dunning, and Joachim I. Krueger, 85–106. New York: Psychology Press.

Anderson, Cameron, Sebastien Brion, Don A. Moore, and Jessica A. Kennedy. 2012. "A Status-Enhancement Account of Overconfidence." *Journal of Personality and Social Psychology* 103(4): 718–35.

Babcock, Linda, and George Loewenstein. 1997. "Explaining Bargaining Impasse: The Role of Self-Serving Biases." *Journal of Economic Perspectives* 11(1): 109–26.

Beauregard, Keith S., and David Dunning. 1998. "Turning Up the Contrast: Self-Enhancement Motives Prompt Egocentric Contrast Effects in Social Judgments." *Journal of Personality and Social Psychology* 74(3): 606–621.

Camerer, Colin F., and Robin M. Hogarth. 1999. "The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework." *Journal of Risk and Uncertainty* 19(1–3): 7–42.

Ditto, Peter H., and David F. Lopez. 1992. "Motivated Skepticism: Use of Differential Decision Criteria for Preferred and Nonpreferred Conclusions." *Journal of Personality and Social Psychology* 63(4): 568–84.

Dunning, David, Judith A. Meyerowitz, and Amy D. Holzberg. 1989. "Ambiguity and Self-Evaluation: The Role of Idiosyncratic Trait Definitions in Self-Serving Assessments of Others." *Journal of Personality and Social Psychology* 57(6): 1082–90.

Festinger, Leon. 1950. "Informal Social Communication." *Psychological Review* 57(5): 271–82.

Festinger, Leon. 1957. A Theory of Cognitive Dissonance. Stanford, CA: Stanford University Press.

Gigerenzer, Gerd. 2004. "Dread Risk, September 11, and Fatal Traffic Accidents." *Psychological Science* 15(4): 286–87.

Gilovich, Thomas. 1991. *How We Know What Isn't So: The Fallibility of Human Reason in Everyday Life.* New York, NY: Free Press.

Gilovich, Thomas, and Lee Ross. 2015. *The Wisest in the Room: How You Can Benefit from Social Psychology's Most Powerful Insights.* New York, NY: Free Press.

Gohmann, Johanna. 2015. "Jimmy Stewart Was My Teen Idol." Salon, December 24. http://www.salon.com/2015/12/24/jimmy_stewart_was_my_teen_idol/.

James, William. 1890. *Principles of Psychology*, vol. 2. New York, NY: Cosimo.

Kahneman, Daniel, and Amos Tversky (eds.) 2000. *Choices, Values, and Frames.* New York, NY: Cambridge University Press and the Russell Sage Foundation.

Klar, Yechiel, and Eilath E. Giladi. 1997. "No One in My Group Can Be Below the Group's Average: A Robust Positivity Bias in Favor of Anonymous Peers." *Journal of Personality and Social Psychology* 73(5): 885–901.

Kunda, Ziva. 1990. "The Case for Motivated Reasoning." *Psychological Bulletin* 108(3): 480–98.

Maass, Anne, Daniela Salvi, Luciano Arcuri, and Gún R. Semin. 1989. "Language Use in Intergroup Contexts: The Linguistic Intergroup Bias." *Journal of Personality and Social Psychology* 57(6): 981–93.

Oster, Emily, Ira Shoulson, and E. Ray Dorsey. 2013. "Limited Life Expectancy, Human Capital, and Health Investments." *American Economic Review* 103(5): 1977–2002.

Pronin, Emily, Thomas Gilovich, and Less Ross. 2004. "Objectivity in the Eye of the Beholder: Divergent Perceptions of Bias in Self versus Others." *Psychological Review* 111(3): 781–99.

Schelling, Thomas C. 1978. *Micromotives and Macrobehavior.* New York, NY: W.W. Norton.

Simon, Herbert A. 1956. "Rational Choice and the Structure of the Environment." *Psychological Review* 63(2): 129–138.

Thaler, Richard H. 1991. *Quasi-Rational Economics.* New York: Russell Sage Foundation.

Trope, Yaacpv, and Nira Liberman. 2003. "Temporal Construal." *Psychological Review* 110(3): 403–421.

Vallacher, Robin R., and Wegner, Daniel M. 1987. "What Do People Think They're Doing? Action Identification and Human Behavior." *Psychological Review* 94(1): 3–15.