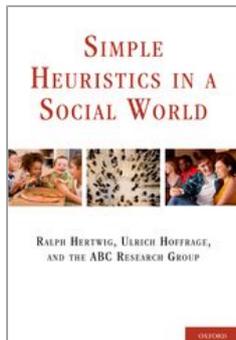


Probabilistic Persuasion: A Brunswikian Theory of Argumentation

University Press Scholarship Online

Oxford Scholarship Online



Simple Heuristics in a Social World

Ralph Hertwig, Ulrich Hoffrage, and ABC Research Group

Print publication date: 2012

Print ISBN-13: 9780195388435

Published to Oxford Scholarship Online: January 2013

DOI: 10.1093/acprof:oso/9780195388435.001.0001

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

Torsten Reimer

Ralph Hertwig

Sanja Sipek

DOI:10.1093/acprof:oso/9780195388435.003.0004

Abstract and Keywords

The Brunswikian lens model has been widely used to describe how individuals integrate information when making a decision (Brunswik, 1943; Dhimi, Hertwig, & Hoffrage, 2004). The chapter applies and extends the lens model to a persuasion context. Specifically, the chapter introduces the *probabilistic persuasion theory* (PPT) as a framework within which the quality of arguments can be defined and measured, and the cognitive processes involved in the selection and in the reception of arguments can be modeled. Construing persuasion within the framework of PPT has the surplus value of opening the door to a rich literature on information processing models in judgment and decision making. The

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

chapter outlines basic assumptions of the new theory, exemplify its application, and discuss its heuristic value. The chapter begins by briefly reviewing dual-process models of persuasion and how they account for the impact of arguments on attitudes. Second, the chapter critically discusses the theories' implications for human rationality, particularly their equation of heuristic processing with irrationality. Third, the chapter describes basic tenets of PPT as an alternative account of persuasion that is based on a Brunswikian framework (Hammond & Stewart, 2001). PPT asserts that persuasion can be construed as a decision-making process, in which a communicator provides information with the goal to influence a receiver's judgments and decisions. The chapter demonstrates how PPT can be used to specify these influence processes and to study the cognitive processes involved in the selection and reception of arguments. Fourth, the chapter derives five testable predictions of the new theory and describe preliminary experimental evidence in support of this account.

Keywords: persuasion, decision making, brunswikian lens model, communication, argument quality, fast and frugal heuristics, social influence, bounded rationality

[B]revity is the best recommendation of a speech, not only in the case of a senator, but in that too of an orator.

Marcus Tullius Cicero (51 B.C./1853)

In the U.S. presidential race of 1960, the Democratic Party nominee, John F. Kennedy, won the general election by a tiny margin. Of nearly 69 million votes cast, only slightly more than 100,000 more votes went to Kennedy than to Richard Nixon, the Republican Party nominee and, as President Eisenhower's Vice President, the quasi-incumbent. In the eyes of many political pundits during the campaign, it was Nixon's election to lose. With the benefit of hindsight, the Kennedy-Nixon debates, which as the first presidential debates to be televised attracted enormous publicity, are now widely seen as the turning point. The first debate, which focused on domestic issues, featured an exchange of views that has a familiar ring even today. As reported in the *New York Times*:

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

Mr. Nixon charged that the Democratic domestic program advanced by Senator Kennedy would cost the taxpayer from \$13,200,000,000 to \$18,000,000,000. […]

“That,” declared Senator Kennedy, in one of the evening's few shows of incipient heat, “is wholly wrong, wholly in error.” […] “I don't believe in big government, but I believe in effective government,” Mr. Kennedy said. (Baker, 1960)

(p.104) Not only the substance of the debate but also the appearance and demeanor of the candidates drew attention and comment. The *New York Times* coverage, for instance, observed:

Senator Kennedy, using no television makeup, rarely smiled during the hour and maintained an expression of gravity suitable for a candidate for the highest office in the land. Mr. Nixon, wearing pancake makeup to cover his dark beard, smiled more frequently as he made his points and dabbed frequently at the perspiration that beaded out on his chin. (Baker, 1960)

The candidates' arguments and demeanor influenced the audience differently, depending on the channel of communication. The majority of people who followed the debate on the radio thought that Nixon won it on substance, whereas most of the 70 million who watched it on television declared Kennedy the winner (see <http://www.museum.tv/>). Nixon learned his lesson. After losing to Kennedy in 1960, he ran for the presidency again in 1968 and for reelection in 1972, but he refused to take part in any more presidential debates, even turning down an offer by his Democratic challenger in 1972, Senator George McGovern, to pay for a nationally televised debate (Kovach, 1972). Nixon handily defeated McGovern in the election that November.

The twentieth century boasts numerous examples of powerful political oratory. Among the most significant American political speeches of that time are Martin Luther King's “I have a dream” speech, John F. Kennedy's inaugural address and “Ich bin ein Berliner” speech, Franklin D. Roosevelt's Pearl Harbor address to the nation, and Malcolm X's “The ballot or the bullet” speech (Lucas & Medhurst, 2008). Since

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

the time of the ancient Greeks, the study of rhetoric had been the preserve of philosophers and historians. Possibly inspired by the twentieth century's great oratory—as well as its corrosive demagoguery—social scientists began in the 1940s and 1950s to investigate the processes underlying rhetoric and persuasion empirically (see Perloff, 2003, for a historical review).

One early finding of this research—disconcerting but, from the perspective of scholars of classical rhetoric, not astounding—was what were assumed to be better arguments do not invariably carry the day (e.g., Petty & Cacioppo, 1986). As the Kennedy–Nixon and other presidential debates demonstrated, voters' opinions can be influenced by many other factors, including the candidates' perceived or actual personality traits and demeanor (e.g., Nixon's five o'clock shadow appears to have projected a sinister image to the television audience); their past (e.g., war hero vs. draft dodger); their experience, maturity, integrity, and competence; and their positions on moral “litmus test” issues (e.g., pro-life vs. pro-choice positions (p.105) on abortion; see Jamieson, 1996). The limits of the power of argument (Kennedy, 1991) to influence opinion gave rise to various psychological models of persuasion.

Among the most influential of these psychological accounts are the *heuristic-systematic model* (Chaiken, 1987) and the *elaboration-likelihood model* (Petty & Cacioppo, 1986). We begin by briefly reviewing both models and how they account for the impact of arguments on attitudes. Second, we critically discuss the theories' implications for human rationality, particularly their equation of heuristic processing with irrationality. Third, we put forth an alternative account of persuasion based on a Brunswikian framework (Hammond & Stewart, 2001). Finally, we describe experimental evidence in support of this account.

Two disclaimers are in order at the outset: We do not consider another influential psychological tradition in persuasion research, one that is more generally concerned with social influence strategies (see Cialdini, 2001). Furthermore, we focus on key common aspects of the heuristic-systematic

model (Chaiken, 1987) and the elaboration-likelihood model, rather than, for instance, comprehensively describing all seven postulates of the elaboration-likelihood model. Let us now turn to the logic behind the two dominant psychological models of persuasion.

Two Cognitive Tools to Evaluate the Speaker's Message

Aristotle distinguished among three means of persuasion. A speech can persuade through the character of the speaker (*ethos*)¹, the emotional state of the listeners (*pathos*), or the argument itself (*logos*; Rapp, 2010). Psychological theories of persuasion are mostly concerned with listeners; their focus, however, is not listeners' emotional state but the information-processing tools that listeners bring to the task of evaluating the speaker's message. The heuristic-systematic model distinguishes between *systematic* and *heuristic* information processing (Chaiken, 1987), whereas the elaboration-likelihood model distinguishes between the *central* and the *peripheral* information-processing routes (Petty & Cacioppo, 1986). Despite the different terminologies, these dichotomies map onto each other (systematic corresponding to the central route, heuristic to the peripheral route), and their respective explanatory successes and limits greatly overlap (for recent expositions of the models, (p.106) see Kruglanski, Erb, Pierro, Mannetti, & Chun, 2006; Petty, Rucker, Bizer, & Cacioppo, 2004).

What separates the two modes of information processing is the cognitive effort that the listener invests to process a message. In particular, both systematic processing and the central route are effortful, whereas heuristic processing and the peripheral route are effortless. Attending to the speaker's credibility or expertise—or what, broadly construed, Aristotle would call the speaker's *character*—means taking account of peripheral cues. Moreover, employing a heuristic such as “Trust this speaker because she is an expert on the subject” would epitomize low-effort, and thus heuristic processing (e.g., Bohner, Ruder, & Erb, 2002; Chaiken, 1987; Reimer, Mata, Katsikopoulos, & Opwis, 2005). Heuristic processing and peripheral cues can be

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

sufficient to decide whether or not to accept a message. In this view, the argument and its quality will carry persuasive weight only if the listener dignifies it with systematic, effortful processing.

What triggers the investment of cognitive effort in evaluating a speaker's message? Empirical investigations suggest that the two key factors are the listener's motivation and ability (e.g., available cognitive capacity; for reviews, see Booth-Butterfield & Welbourne, 2002; Todorov, Chaiken, & Henderson, 2002). If a listener is highly motivated and able to scrutinize a message, processing will be systematic. If, however, a listener lacks the motivation or the capacity to scrutinize a message, processing will be doomed to be heuristic (Petty et al., 2004).

A 1981 study by Petty, Cacioppo, and Goldman—a classic investigation in the tradition of dual-process models of persuasion—illustrates how both the elaboration-likelihood model and the heuristic-systematic model have typically been tested (see Figure 4-1). The experimenters asked undergraduate students to listen to an audiotaped message about purported changes in the university's graduation requirements. According to the message, all undergraduates would be required to take senior comprehensive exams in order to graduate. The participants' attitude—a common target variable in this area of research—toward such comprehensive exams was the dependent measure. Half of them were told that the new policy would be implemented in one year (high involvement), whereas the other half learned that the new policy would be implemented in ten years (low involvement). In addition, the message was attributed either to the Carnegie Commission on Higher Education (peripheral cue: high expertise) or to the local high school (low expertise). Finally, the policy change was supported by arguments of either high or low quality. In this and many similar experiments, the results are interpreted as follows: If the peripheral cue (in this case, expertise) affects the listener's attitude, it is inferred that the message's processing was heuristic. Conversely, (p.107)

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

if argument quality shapes the listener's attitude, it is inferred that the message was processed systematically (Figure 4-1).

That is, the mode of processing is *inferred* from the effects attributed to

cues and arguments, respectively. In other words, involvement is assumed to trigger the mode of processing, which, in turn, amplifies or attenuates the impact of expertise and argument quality.

Petty et al. (1981) found that when students' involvement was low, their attitudes were influenced mostly by the expertise cue. This effect was interpreted as conforming to the assumption that low involvement triggers heuristic information processing. When the policy change had the potential to affect students directly, in contrast, their attitudes were shaped only by the arguments' quality. This effect was interpreted as conforming to the assumption that high involvement triggers systematic information processing. The established conclusion from these and similar findings is that good arguments sway listeners' attitudes or judgments only when listeners are not on "autopilot" but instead devote their mental capacities to systematically poring over the arguments. Conversely, attributes such as the speaker's expertise are assumed to shape listeners' attitudes when they fail to subject the arguments to more than heuristic processing (e.g., Chen & Chaiken, 1999; Petty et al., 2004).

Dual-Process Models: Vague Dichotomies and the Separation of Rationality

Dual-process models have been successfully employed across a variety of persuasion and communication contexts (e.g., Chen & Chaiken, 1999; Petty et al., 2004). At the same time,

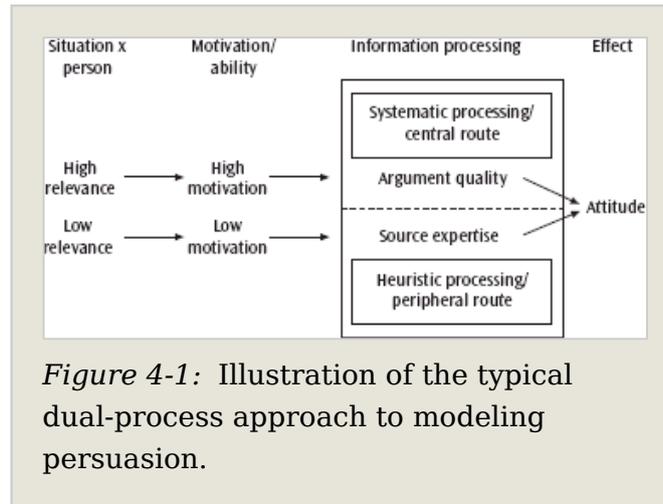


Figure 4-1: Illustration of the typical dual-process approach to modeling persuasion.

they have met with vigorous criticism (e.g., Hamilton, Hunter, & Boster, 1993; Mongeau & Stiff, 1993; Stiff, 1986). In what follows, we are not concerned with the models' empirical record (e.g., Johnson & Eagly, 1989) or (p.108) with possible experimental confounds (e.g., Pierro, Mannetti, Erb, Spiegel, & Kruglanski, 2005) but with three conceptual issues.

What Is Behind the Labels?

Two dichotomies underpin dual-process models of persuasion. The first is that between heuristic and systematic processing; the second, between cues and arguments. Challenging both dichotomies, Kruglanski and Thompson (1999a, 1999b) argued that peripheral cues (e.g., expertise, credibility) and arguments are functionally equivalent; that is, that cues can take the role of arguments. Moreover, if the two are inseparable, then by extension their assumed modes of processing will be inseparable as well. Proponents of dual-process models of persuasion would be in a position to counter this conclusion if the conjectured processes were measured independently. As emphasized earlier, however, the mode of processing is commonly inferred from effects attributed either to cues or to arguments, respectively. Taking aim at this inferential leap, Stiff (1986) wrote that the elaboration-likelihood model is a "model of human information processing centering on the strategies individuals use to process information. However, Petty and Cacioppo fail to assess directly the cognitive processes themselves" (p. 77).

The ultimate reason why the cognitive processes hypothesized to underlie persuasion have not been directly captured may be that they tend to be "one-word" explanations; that is, explanations in which a word (e.g., *systematic*, *heuristic*), usually broad in meaning, is the explanans. However apt a description, the word does not specify an underlying mechanism or a theoretical structure, and thus can hardly constrain researchers in their use of it (Gigerenzer, 1998, p. 2). For instance, where dual-process proponents see the influence of a speaker's expertise squarely as a reflection of low motivation and reliance on heuristic processing, others have argued that Petty et al.'s (1981) findings are consistent with the view that the expertise of the message's source can

affect listeners even when they are highly motivated. In this second view, high relative to low motivation may simply alter how expertise is inferred rather than the operation of an underlying trust-the-expert heuristic; that is, among highly motivated listeners, whether a speaker is perceived to be an expert may depend on the merits of his arguments (Reimer, Mata, & Stoecklin, 2004; Reimer et al., 2005).

What Makes an Argument Good?

Dual-process models of persuasion typically pit arguments against peripheral cues and attribute superior quality to arguments. This (p.109) attribution rests on a purely empirical foundation: Argument quality is validated through the subjective judgments of respondents. Consequently, dual-process models lack a theoretically rooted criterion for the quality of argumentation. More generally, scholars of communication science (O'Keefe, 2003; Stiff, 1986) have bemoaned that experimental research on persuasion lacks a theoretical definition of what makes an argument “good”: logical coherence? simplicity? accuracy? a combination of these? Or is it something else altogether? Without a theory of the quality of arguments—and of cues—it is impossible, for instance, to exclude the possibility that people heed peripheral cues simply because they consider them to be worthier than the presented arguments.

Why Should Heuristic Processing Be Irrational?

Dual-process models of persuasion rest on a popular distinction in research on social cognition and cognitive psychology that splits the mind into two qualitatively different processes or systems. Dual-process models, of which there are many, presuppose that heuristic (intuitive) and systematic (deliberate) processes are aligned with certain properties. Heuristic processing has been portrayed as associative, quick, unconscious, effortless, heuristic, and, importantly, error-prone. Systematic processing, in contrast, has been depicted as rule-based, slow, conscious, effortful, analytical and, importantly, rational. Conjectures about the existence of two separate processing systems have been buttressed by abundant empirical findings that have been interpreted to support the duality of the mind (e.g., for reviews, see Evans, 2008; Kruglanski & Orehek, 2007). At the same time, the

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

dualistic view of human cognition and its implications for rationality have also been incisively criticized (e.g., Keren & Schul, 2009; Kruglanski & Gigerenzer, 2011).

One key point of criticism concerns the equation of heuristic processing and suboptimal performance. The article of faith behind this equation is that the more laborious, computationally expensive, and nonheuristic the cognitive strategy, the better the judgments to which it gives rise. This view reflects a conception of heuristics that emerged in research on social cognition and decision making in the 1970s as overused, mostly dispensable cognitive processes that people often apply to situations where rules of logic and probability theory should be used instead (e.g., Gilovich, Griffin, & Kahneman, 2002; Kahneman, Slovic, & Tversky, 1982). Heuristics were thus fingered as the cognitive culprits behind an extensive catalogue of violations of norms taken from probability theory, logic, and statistics. Why do people resort to using such third-rate cognitive software? The typical answers to this question have been that people use heuristics either (p.110) because of the former's cognitive limitations or to save effort at the expense of accuracy. The first reason implies an inability to optimize and perform rational calculations; the second reason implies a pragmatic decision that doing so may not be worthwhile. Both rest on a principle that is often taken to be a general law of cognition; namely, the accuracy-effort tradeoff. The less information, computation, or time that one uses, the less accurate one's judgments will be (see Gigerenzer, Hertwig, & Pachur, 2011).

A different view of heuristics has been laid out by Gigerenzer, Todd, and the ABC Research Group (1999), Todd, Gigerenzer, and the ABC Research Group (2012), and the authors of this volume. Inspired by Herbert Simon's (1990a) concept of bounded rationality, this view holds that the human "cognitive toolbox" includes heuristics because their building blocks—for instance, limited search, stopping rules, one-reason decision making, and aspiration levels—can lead to more accurate inferences or predictions than can algorithms based on the principles of logic, probability, or maximization (e.g., Gigerenzer & Brighton, 2009). Thus, depending on a

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

heuristic's ecological rationality (the degree to which it is adapted to the structure of an environment), less effort can lead to higher accuracy (chapter 1). One key to the success of heuristics is their robustness; that is, their ability to operate successfully when the environment changes. Robustness often follows from simplicity—the signature of a heuristic—because simple models with few or no free parameters are less vulnerable to overfitting (i.e., increasing the model fit by accommodating noise: see Gigerenzer et al., 2011). Although the view that heuristics reflect inferior reasoning is still widespread in research on social cognition and social perception, some researchers in this area have underscored that heuristics can be surprisingly accurate when used in appropriate social environments (Funder, 1987; McArthur & Baron, 1983; Swann, 1984).

To conclude, arguments do not unfold in a pristine sphere of ideas. Instead, they compete in a marketplace in which myriad factors beyond an argument's intrinsic quality—for instance, Kennedy's vaunted charisma and Nixon's less than telegenic demeanor—determine whether an argument holds sway. Classic psychological theories of persuasion attribute the impact of factors other than issue-relevant arguments to a heuristic processing style that is assumed to be suboptimal. Argument quality prevails only when people bother to invest sufficient effort to scrutinize the message. This dual-process view has been criticized for its lack of specified processes (despite the emphasis on modes of information processing) and a theoretical benchmark for argument quality, as well as for its frequent equation of heuristic processing with faulty cognitive software.

Not least because of the criticisms just mentioned, Kruglanski and Thompson (1999a, 1999b) proposed a *unimodel* of persuasion (p.111) that puts peripheral cues (e.g., expertise) on a par with arguments as potential *evidence* for a standpoint. The extent to which evidence affects a listener's judgment depends on several dimensions, including perceived task difficulty, processing motivation, cognitive capacity and motivational biases, and the order in which evidence is presented and processed (Erb et al., 2003). The unimodel is a

parametric model. It represents the postulated dimensions in terms of parameters and, depending on the parameter values, predicts different persuasive effects on the listener. For example, if the task difficulty is perceived to be high, evidence is expected to have an effect only on listeners with sufficient processing capacity.

In what follows, we propose a new theoretical framework of persuasion. Inspired by Kruglanski and Thompson's approach (1999a, 1999b), it is built on the assumed functional equivalence of peripheral cues and arguments. It also shares the unification view laid out by Kruglanski and Gigerenzer (2011), according to which both systematic and heuristic processing are based on rules; that is, inferential devices that can be described in terms of “if-then” relations of the type “if (cues), then (judgment).” Our framework differs from Kruglanski and Thompson's unimodel in that it rests on Egon Brunswik's (1952) *probabilistic functionalism* and an interpretation of the Brunswikian lens model based on simple heuristics (see Gigerenzer & Kurz, 2001)—building blocks that we explain in detail shortly. Most important, departing from the premise in dual-process models that heuristics constitute suboptimal shortcuts to normative calculations, we treat heuristics as valuable assets that enable human communication and inference.

Some Boundaries and a Fictitious Presidential Debate

Let us first be clear about the many things our framework cannot accommodate. In order to define the boundaries, some time-honored distinctions can help. Of the three means of persuasion described by Aristotle (see Kennedy, 1991; Rapp, 2010), we are concerned with the character of the speaker (in terms of, say, expertise and credibility) and the argument itself, but not with the emotional state of the listener. Aristotle also identified three “species” of rhetoric. *Deliberative* and *judicial* speech, which takes place in the assembly or before a court, puts the listener in the position of having to decide in favor of one of opposing parties, standpoints, or actions. *Epideictic* speech, in contrast, praises or blames somebody.

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

Finally, Aristotle distinguished between two kinds of arguments: inductions and deductions. Induction is defined as an argument that proceeds from the particular to a universal, whereas a deduction is an (p.112) argument in which, given certain premises, something different necessarily arises from the premises. Our focus here is on deliberative and judicial speech and on messages involving inductive arguments (but let us also emphasize that the distinction between induction and deduction is likely to be obsolete in explanations of human reasoning; see Oaksford & Chater, 1996). Furthermore, we assume that the speaker does not intentionally deceive the listener and that the listener strives to hold accurate views of the world (see Petty & Cacioppo, 1986). Finally, dual-process theories of persuasion have commonly focused on attitudes. Like Gonzalez-Vallejo and Reid (2006), we believe that successful persuasion must ultimately manifest itself in behavioral changes. With our probabilistic persuasion theory (PPT) and its focus on choice and judgment, we hope to get closer to behavior.

With these boundaries in mind, we now turn to a purely fictitious exchange of arguments that we will use henceforth to illustrate the present framework. The context of the exchange is that of a televised American presidential debate on domestic policy between the Republican and Democratic presidential nominees. The nominees' target of persuasion is the debate's television audience. The moderator's first question concerns the pressing problem of homelessness in U.S. metropolitan areas:

Moderator:

Welcome. Let's get to it. A recent article in the *New York Times* painted the following bleak picture: Dozens of U.S. cities across the country deal "with an unhappy déjà vu: the arrival of modern-day Hooverilles, illegal encampments of homeless people that are reminiscent, if on a far smaller scale, of Depression-era shantytowns" (McKinley, 2009). Moreover, *The Economist* recently reported the heart-wrenching fact that "during the 2008–2009 school year, America's public schools reported more than 956,000 homeless pupils, a 20% increase over the previous school year" ("Getting Strategic," 2010). Let me make the

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

homelessness crisis in our cities as concrete as possible. Governor, you grew up in Phoenix; the Senator is from Boston. Do you have any idea which of the two cities suffers from more homelessness? Governor, you go first, please.

Governor:

First of all, let me say that it is not acceptable for children and families to be without a roof over their heads in a country as wealthy as ours. Second, let me admit that I do not know the exact numbers for Boston and Phoenix. But I do know that urban planners and economists have identified numerous factors that predict homelessness, including rent control, average temperature, unemployment rate, housing vacancy rate, and the proportion of people living below the poverty line. To the best of my knowledge, the most powerful predictor is average temperature. In all likelihood, Phoenix is bound to have a higher rate of homeless people than Boston. It's simply warmer there, and there is little the government can do about our climate.

Moderator:

All right, thank you. Senator?

(p.113) *Senator:*

My impression is that the governor just let slip us how little he is willing to do about global warming—but never mind that for now. I think everybody understands at this point that a few years ago we experienced the worst financial crisis since the Great Depression. The senator and I agree that it is not acceptable for American families to be without a roof over their heads. I disagree with the governor, however, and …

Moderator:

Senator, allow me to interrupt and simply ask: Do you know whether your home city has more or fewer homeless people than the governor's?

Senator:

Well, I don't, but like the governor I am aware of the opinions of economists and urban planners. To the best of my knowledge, the best predictor of homelessness is rent control. Why? In my view, rent regulations, despite good intentions, prevent housing creation, raise prices, and increase urban blight. Now, I happen to know that Phoenix has abolished rent control, while my hometown, Boston, has

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

kept it. So I disagree with the governor. To my chagrin, I believe that Bostonians these days are experiencing a higher rate of homelessness in their streets than are the residents of Phoenix. And unlike the senator, I believe there is something we can do about it!

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

Our probabilistic persuasion theory (PPT) rests on two pillars: the *Brunswikian lens model* (Brunswik, 1952) and, building on it, the notion of a *fast and frugal lens model* (Gigerenzer & Kurz, 2001). We will explain both in detail. But first, a preview. The lens model allows us to conceptualize listeners' frame of mind and how they process the speaker's message and, equally important, provides us with a criterion for argument quality. To this end, let us replay the fictitious debate in "fast motion."

The moderator assigns the speakers a task in which it must be inferred which of two objects has a higher value on a criterion. Examples of tasks with this structure abound: allocation of financial resources (e.g., which of two education acts should be implemented and funded, with student performance as the criterion); policy decisions (e.g., which of two environmental policies should be enacted, with carbon dioxide emissions as the criterion); and, as in the present case, sociodemographic predictions (e.g., which of two cities has the higher rate of homelessness, crime, or mortality). Tasked by the moderator to judge which of their respective home cities has a worse homelessness crisis, they each admit to being caught on the hop. To compensate for their lack of direct knowledge, they select predictors of homelessness, stress the predictive validity of the selected predictors, and on the basis of them come to opposite conclusions. How can a listener evaluate and process the speakers' messages to determine which one has the better arguments?

(p.114) The Brunswikian Lens Model and Vicarious Functioning

Let us assume that the listener, like the speaker, has no certain knowledge of the cities' homelessness rate; otherwise, she would simply retrieve it. For instance, a person may recall

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

having recently read that Phoenix belongs to the five U.S. cities with the highest rates of homelessness, and that Boston was not in this group. Complemented by elementary logical operations, this knowledge would be sufficient to answer that Phoenix has a higher homelessness rate than Boston and therefore to conclude that the governor's message is accurate. Although such "local mental models" (Gigerenzer, Hoffrage, & Kleinbölting, 1991) provide a neat solution to the task, they are probably used rarely in real-time exchanges where the listener cannot consult external knowledge sources.

If no local mental model can be constructed, the listener can nevertheless intuit the answer by linking up the specific task with the probability structure of a corresponding natural environment. According to Brunswik's (1952) theory of probabilistic functionalism, the environment offers (proximal) cues; that is, variables that covary with the (distal) criterion of interest. The mind's cognitive and perceptual inference machinery can thus take advantage of cues to infer criteria that are not directly observable. The main tenets of Brunswik's probabilistic functionalism are illustrated in his lens model, presented in Figure 4-2.

The double convex lens shows a collection of proximal stimuli (cues) diverging from a distal criterion (or outcome) in the environment. When the distal criterion to be inferred is the distance of an object to the organism, for instance, the cues might be the retinal size of the stimulus object, aerial perspective, occlusion, and retinal disparity (stereopsis). When the distal criterion is a city's homelessness rate, possible cues include rent control, average temperature, unemployment rate, and vacancy rate. Not all these cues are of equal utility. Brunswik (1952) proposed measuring the *ecological validity* of a cue by the Pearson correlation between the cue and the distal variable (Figure 4-2). Validity's counterpart is *utilization*; that is, the degree to which the organism makes use of available cues. With *achievement*, Brunswik described the degree to which perception (or cognition) captures the distal stimulus, measured in terms of the correlation between the distal criterion (e.g., actual distance) and the response of the organism (e.g., estimated distance). The lens model describes the organism and environment as part of the same system, as

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

“equal partners” in a relationship that “has the essential characteristic of a ‘coming-to-terms’ ” (Brunswik, 1957, p. 5).

The environment an organism must adapt to is not perfectly predictable from cues (Brunswik, 1943). For example, a retinal projection of a given size can indicate either a large object that is far away or a

(p.115)
small object that is close. Moreover, a given cue may not always be present. In other words, cues are uncertain indicators of the distal criterion. Therefore, an adaptive system relies on multiple cues that can be substituted for one other because they are interrelated (see the

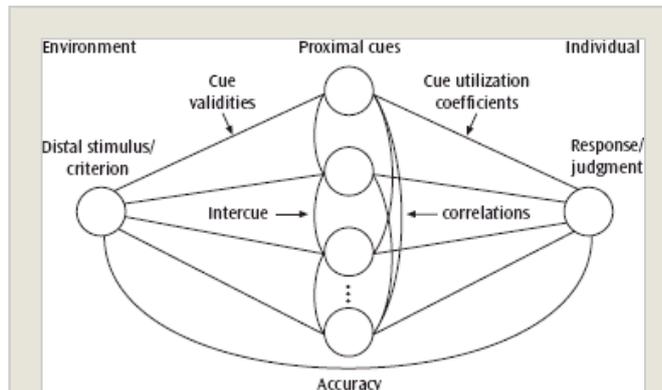


Figure 4-2: Adapted lens model.

(Source: Adapted from Figure 1 in “The role of representative design in an ecological approach to cognition” by M. K. Dhami, R. Hertwig, & U. Hoffrage [2004], *Psychological Bulletin*, 130, 959–988. Copyright 2004 by the American Psychological Association.)

intercue correlations in Figure 4-2). Such flexible cue substitution, known as *vicarious functioning*, has frequently been modeled by multiple regression (see Hammond & Stewart, 2001). This choice, however, has come under criticism, and an alternative model has been proposed.²

A Fast and Frugal Lens Model

Gigerenzer and Kurz (2001) observed that the neo-Brunswikian modeling of vicarious functioning in terms of multiple regression presupposes two fundamental processes; namely, the weighting of cues (by their correlations with the distal criterion) and the summing of cue values. Although weighting and summing have been used to define rational

judgment since the Enlightenment—expected value and expected utility theories, for instance, rest on both processes—they have also been challenged. In particular, the question has been raised of to what (p.116) extent their combination can result in a model of human cognition that respects the limitations of human time and knowledge.

In what follows, we offer a *fast* and *frugal* lens model of vicarious functioning that is intended as an alternative to multiple regression (see Hammond & Stewart, 2001). *Fast* and *frugal* refer to cognitive processes that enable the organism to make inferences under conditions of limited time and information. Unlike multiple regression, a fast and frugal lens model does not aim to integrate all cues into one judgment. Instead, it applies heuristic principles for information search, stopping search, and inference. For processing cues, the take-the-best heuristic (Gigerenzer & Goldstein, 1996), derived from the theory of probabilistic mental models (Gigerenzer et al., 1991), provides a powerful alternative to multiple regression. For simplicity, we assume that all cue values are binary; that is, either positive or negative (with positive values indicating higher homelessness rates in the example above). We also ignore the first step of the take-the-best heuristic, the *recognition heuristic*, which we return to later. The take-the-best heuristic can be expressed in the following steps:

Step 1. *Search rule*: Choose the cue with the highest validity that has not been tried for this choice task.

Look up the cue values of the two objects.

Step 2. *Stopping rule*: If one object has a positive cue value and the other does not (i.e., either negative or unknown value), then stop search and go to Step 3; otherwise return to Step 1 and search for another cue. If no further cue is found, then guess.

Step 3. *Decision rule*: Predict that the object with the positive cue value has the higher value on the criterion.

This fast and frugal lens model relies on one-reason decision making. That is, in contrast to multiple regression, the inference is based solely on the most valid cue that discriminates between the objects. It may be wrong, but none

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

of the remaining cues, nor any combination of them, can change it. In other words, the take-the-best heuristic is a noncompensatory strategy. Its search order is determined by the ranking of cues according to their validities v_i :

$$v_i = \frac{R_i}{R_i + W_i},$$

where R_i is the number of correct inferences and W_i is the number of incorrect inferences based on only one cue i (among all pairs of objects in which the cue discriminates; that is, one object has a positive value and the other does not). Ranking cues according to their validity is relatively simple, as it ignores, among other things, the (p.117) dependencies between cues (which multiple regression takes into account). Although this cue ranking is not “optimal” (Martignon & Hoffrage, 2002), Gigerenzer and Brighton (2009; also Gigerenzer & Goldstein, 1996; Katsikopoulos, Schooler, & Hertwig, 2010) demonstrated that take-the-best, when tested in an environment in which the order of cues was not known but had to be estimated from limited samples, could make more accurate predictions than strategies that use all possible information and computations, including optimization models. Figure 4-3 illustrates a fast and frugal lens model based on the take-the-best heuristic. To avoid misunderstanding, let us emphasize that take-the-best is only one possible manifestation of a fast and frugal lens model of persuasion; other heuristics could easily take the place of take-the-best in our Brunswikian framework. With the fast and frugal lens model in place, we can now explicate PPT. Before we turn to how listeners process arguments, let us first define argument quality using the lens model.

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

(p.118)

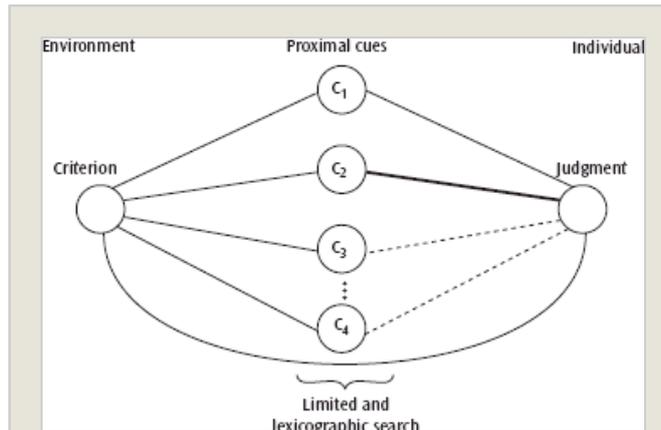
Probabilistic Persuasion Theory: Validity of Arguments (Cues)

The speaker in a deliberative or judicial speech conveys information with the goal of informing and influencing others' choices.

Listeners process and evaluate this information and decide in favor of one of the advocated positions.

Consider, for illustration, the governor in our fictitious debate. Of several mentioned cues, he

selects temperature as the best predictor. The senator, in contrast, selects rent control. Based on those cues, they arrive at different inferences. Which of the two inferences should the listener buy into?



*Figure 4-3: Illustration of a fast and frugal lens model. The task is to infer which of the two objects (e.g., two cities) has a higher value on a quantitative criterion (e.g., homelessness). For simplicity, cues (C₁ to C₄) are assumed to be binary, looked up in the order of their validity. The first cue, C₁, does not discriminate between objects (fine line), but the second one does (thick line). Search is therefore terminated, and the inference is made on the basis of the values of C₂. The cue values of C₃ and lower-ranked cues are not searched (broken lines). (Source: Adapted from Figure 24.1 of "The vicarious function reconsidered: A fast and frugal lens model" by G. Gigerenzer & E. M. Kurz [2001] in K. R. Hammond & T. R. Stewart (eds.), *The essential Brunswik: Beginnings, explications, applications*. New York: Oxford University Press. Copyright by Oxford University Press)*

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

PPT assumes that the answer to this question will depend on the perceived cue validities. Kruglanski and Thompson (1999a, 1999b) argued that cues and arguments are functionally equivalent. Although those authors were concerned with peripheral cues (e.g., credibility), we generalize their premise: Cues of any kind can be put forth as arguments. If so, then cue–argument equivalence implies that argument quality can be derived from the goodness of cues as measured by ecological validity (henceforth we use the terms *cue* and *argument* interchangeably). The fast and frugal lens model (Figure 4-3) defines ecological validity in terms of the relative frequency with which a cue correctly predicts the criterion (see the equation above) in a specific reference class (Brunswik, 1943, p. 257); that is, a specific category of objects or events (in our example, the largest U.S. cities). The reference class determines which cues can function as probability cues for the criterion and what their validities are (Hoffrage & Hertwig, 2006). Ecological validities are thus a measurable indicator of the quality of arguments: The higher a cue's ecological validity, the stronger the respective argument that uses this cue.

Table 4-1 reports the ecological validities of eight cues in predicting homelessness rates in the 50 largest U.S. cities. The validities range widely, with rent control—the predictor emphasized by the senator—being the most valid cue. A person who relies exclusively on this cue to infer which of two U.S. cities has a higher homelessness rate will be correct in 90% of cases (in which the cue discriminates between the two cities to be compared). In contrast, the average-temperature cue—the predictor underscored by the senator—has a validity of 69%. By this measure of argument quality, the senator has the better argument (and, indeed, according to the 2010 Survey of the United States Conference of Mayors, Boston's homelessness rate is likely to be higher than that of Phoenix).

Ecological validities offer researchers of persuasion an objective criterion for defining argument quality (for an alternative, coherence-based approach to defining argument strength, see Pfeifer, 2007; for a Bayesian approach, see Hahn & Oaksford, 2007). Evaluating arguments in terms of the

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

goodness of cue measures also allows for models that go beyond two-alternative choices. Furthermore, cue-goodness

(p.119)

Table 4-1: Cues Predictive of the Homelessness Rates in the 50 Largest U.S. Cities

Cue	Definition	Ecological validity
Rent control	Does the city have rent control?	0.90
Average temperature	What is the city's average temperature?	0.68
Unemployment rate	What is the city's rate of unemployment?	0.59
Population	What is the city's population size?	0.58
Poverty	How many residents' income is below poverty line (in %)?	0.54
Vacancy rate	How many buildings are vacant (in %)?	0.43
Public housing	How many people live in public housing (in %)?	0.41

Notes. The cues to homelessness were taken from Tucker (1987). We updated the cue values and the criterion where possible. Continuous variables were dichotomized on the basis of a median split (see Czerlinski, Gigerenzer, & Goldstein, 1999). Note that, for cues with validities above 0.50, the city with the larger cue value has a higher homelessness rate than the city with the lower cue value in most pairs in which the cue discriminates. Conversely, for cues with validities below 0.50, the city with the higher cue value has a lower homelessness rate than the city with the lower cue value in most pairs in which the cue discriminates.

measures can be defined with regard to other cognitive tasks, such as estimation and classification (Gigerenzer et al., 1999).

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

Experimenters can calculate cue validities using a reference class and cue information. But how well developed is people's intuitive sense of the validity of a cue and, by extension, of an argument? Gigerenzer et al. (1991) assumed that the more experience people have with a reference class, a target variable, and cues in their environment, the more closely their estimates of cue validities will correspond to ecological validities. Relatedly, Katsikopoulos et al. (2010) showed that people have surprisingly good intuitions about the direction of the correlations between cues and criterion. Nevertheless, a listener's subjective cue order will not invariably map onto that of the ecological validities (as, for instance, appears to be the case in the field of deception and lie detection; see DePaulo et al., 2003; Levine, Kim, Park, & Hughes, 2006; Sporer & Schwandt, 2007). Yet any subjective cue order will endow the listener with a benchmark for judging the (subjective) quality of arguments.

Finally, not all cues are created equal. PPT proposes distinguishing among four categories. Cues can reflect objective properties of an object; objective or perceived properties of the speaker and the context, respectively; or the knowledge state of the listener. Table 4-2 lists these four cue categories and illustrations thereof. Regardless of their classification, all these cues have predictive power that can (p.120) be quantified in terms of cue validity or another measure of cue goodness.

Probabilistic Persuasion Theory: How the Listener Processes the Speaker's Arguments

How does the listener process and evaluate the speaker's arguments? Processing in verbal communication must be fast. Within just a few moments, the listener needs to grasp the meaning of referents, retrieve relevant knowledge from memory, and at least implicitly evaluate the quality of the arguments. A review of the abundant research in linguistics and psycholinguistics on what makes such rapid processing possible lies beyond the scope of this chapter. Interestingly, however, the processes underlying verbal comprehension have been described as tantamount to the lexicographic processing of fast and frugal heuristics. According to Wilson and Sperber, for instance, to communicate is to claim someone's attention

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

and thereby to imply that the information communicated is relevant:

The relevance-theoretic comprehension procedure … (“Follow a path of least effort in computing cognitive effects; test interpretive hypotheses in order of accessibility; stop when your expectations of relevance are satisfied”) could be seen as a “fast and frugal heuristic,” which automatically computes a hypothesis about the speaker's meaning on the basis of the linguistic and other evidence provided. (2004, p. 625)

Indeed, substantial experimental work has demonstrated that simple heuristics are most likely to be used when time is short and information has a cost (e.g., needs to be retrieved from memory, see Gigerenzer et al., 2011). These conditions are typical of verbal communication.

PPT assumes that the listener's default processing of arguments is in terms of fast and frugal heuristics such as take-the-best. Importantly, it distinguishes between listeners with and without cue knowledge. Listeners without knowledge of cues and cue validities (including recognition or fluency knowledge; Goldstein & Gigerenzer, 2002; Hertwig, Herzog, Schooler, & Reimer, 2008) in the respective domains will, *ceteris paribus*, take the cues and cue values embedded in a speaker's message at face value and process them via a simple heuristic. Such listeners will thus not enrich the mental model of the task constructed by the speaker with their own cue knowledge about the objects (see Table 4-2). If, however, two speakers contradict each other—as is the case in our fictitious debate—listeners may resolve the conflict by taking into account speaker cues (e.g., demeanor, perceived expertise, credibility, and party affiliation; Table 4-2). (p.121)

Table 4-2: Four Categories of Cues and Examples Thereof

Category	Example
<i>Object cues</i> : Cues reflecting objective properties of objects	The average temperature of cities

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

Category	Example
<i>Speaker cues</i> : Cues embodying objective or perceived properties of speakers	Objective: e.g., speaker's gender Perceived: e.g., speaker's demeanor*
<i>Listener cues</i> : Cues embodying knowledge about the object that is specific to the listener	Recognition (Goldstein & Gigerenzer, 2002) and fluency (Hertwig, Herzog, et al., 2008)
<i>Context cues</i> : Cues embodying objective properties of the conversational context	Other listeners' response to the message (e.g., heckling or applause)

(*) Depending on the circumstances, a speaker cue such as expertise could be either objective (e.g., the speaker's academic credentials) or subjective (e.g., the speaker's perceived confidence).

In contrast, a listener with cue knowledge faces a choice. He can choose to focus on the cue knowledge included in the speakers' messages and process it. For example, he could evaluate the speakers' conclusions regarding the relative rates of homelessness in Phoenix and Boston by employing the take-the-best heuristic to process the two cues selected by the speakers; namely, average temperature and rent control. In that case, he need only determine which of the two cues ranks higher (with respect to perceived validity) in order to decide which speaker's conclusion he endorses. Alternatively, he can go beyond the given information and bring new cues to the task. For instance, if the listener happens to know of a cue that exceeds the validity of the cues identified by the governor and senator, and he happens to know Phoenix's and Boston's values on this cue, he will be able to exploit this cue. Which strategy he chooses to pursue—focusing on the information given or going beyond it—depends on various factors, such as time pressure, whether or not the speakers' arguments conflict with each another, his perception of the speakers' credibility and expertise, and his confidence in his own knowledge (perceived self-expertise).

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

Probabilistic Persuasion Theory: Means of Persuasion

PPT assumes that a listener evaluates arguments (cues) according to their cue validity (or other measures of goodness) and processes them by employing a noncompensatory heuristic such as take-the-best. What routes can the speaker take in order to persuade the listener? In principle, there are four. The first three concern the listener's knowledge: The speaker can aim to mold the listener's cue knowledge by embedding in the message specific cues (e.g., “rent control”); specific cue values (of objects on cues; e.g., “Boston has rent control but Phoenix does not”); and the validity of specific cues (“rent control is the best cue”). Taking the fourth route, the speaker can target how the listener processes this cue knowledge by suggesting a strategy that differs from noncompensatory processing. For instance, the speaker can list numerous arguments and appeal to the listener to take all arguments, independent of their validity, into account. If successful, such an appeal could prompt the listener to apply a simple compensatory *tallying* heuristic rather than take-the-best. Dispensing with the weighting of arguments according to their quality, this compensatory heuristic simply sums the cue values:

Step 1. *Search rule*: Search through all cues in random order. Look up the cue values.

Step 2. *Stopping rule*: After m ($1 < m \leq M$) cues, stop search and determine which object has more positive cue values, and go to Step 3. If the two tallies are equal, return to Step 1 and search for another cue. If no more cues are found, go to Step 3.

Step 3. *Decision rule*: Predict that the object with the higher number of positive cue values has the higher value on the criterion. If the objects tie with respect to this number, guess.

Different versions of the tallying heuristic exist: some assuming that all ($m = M$) and others that only m significant cues are looked up (Dawes, 1979).

In sum, based on the premise of cue-argument equivalence, PPT employs ecological validity (or related measures of cue goodness) as an objective benchmark for argument quality. It

also proposes that the listener's default style of processing verbal arguments (cues) is noncompensatory and can be modeled in terms of fast and frugal heuristics (such as take-the-best; Figure 4-3). Finally, the theory delineates how a speaker can persuade the listener by molding the listener's cue knowledge or by altering the default processing strategy (e.g., by shifting the processing from noncompensatory to compensatory). We now turn to five predictions derived from our Brunswikian framework.

Probabilistic Persuasion Theory: Predictions

The predictions of PPT concern the selection of arguments by the speaker, the role of speaker and listener cues, the difference between verbal and written messages, the impact of the listener's state of cue knowledge, and the match between the speaker's and listener's respective lens models. We first state each prediction and then explain the rationale behind it.

(p.123) *Prediction 1. Preference for a few good arguments: Speakers select arguments (cues) according to some measure of goodness and tend to focus on a few good ones rather than presenting all arguments available.*

How do speakers decide which arguments to embed in their messages? Dual-process models of persuasion do not specify how arguments are generated and selected, but they do suggest that speakers who want to be persuasive ought to use as many arguments and of as high a quality as possible. In contrast, PPT suggests that the speaker's selection of arguments is guided by measures of argument goodness and that by no means will all available arguments be included in the message. This prediction is derived from two premises. First, according to Grice (1975, 1989), conversations are to some degree cooperative and coordinated efforts (Clark, 1996a), and participants are therefore expected to observe specific *maxims*. Several Gricean maxims exhort the speaker to focus on a few good reasons—for example, “Do not make your contribution more informative than required,” “Be relevant,” and “Be brief.” Second, according to PPT, listeners need to make inferences under conditions where time is

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

limited and information has a cost (e.g., requires memory retrieval). Speakers in online (not scripted) communication find themselves operating under the very same conditions. Therefore, speakers will tend to insert a few good arguments into the message rather presenting all available arguments.

Which arguments are “good”? PPT suggests different measures of goodness depending on the cognitive task to be performed. Taking ecological validity as a definition of argument goodness in tasks requiring a decision about which of two options, events, or objects scores higher on a criterion (e.g., homelessness), a speaker is predicted, *ceteris paribus*, to be more likely to embed a cue with high validity in his message than a cue with low validity. As preferential choices lack an objective criterion, modeling tasks of this type requires replacing ecological validity with another measure of goodness, such as cue distinctiveness—an issue we will return to shortly.

Prediction 2. Primacy of cognitively inexpensive cues: Arguments in terms of object cues can be overruled by speaker and listener cues even when listeners are highly motivated and processing capacity is not compromised.

This prediction builds on the distinction between different categories of cues (Table 4-2). Like object cues (e.g., rent control), speaker cues such as expertise, likeability, and credibility, and listener cues such as recognition and fluency have a quantifiable predictive (p.124) potential (see e.g., Goldstein & Gigerenzer's notion of recognition validity [2002]; and Hertwig, Herzog, et al.'s notion of fluency validity [2008]). Take the expertise cue as an example. Let us assume that in our fictitious debate there is an independent presidential candidate who currently leads the U.S. Department of Health and Human Services. This participant's longtime field of policy expertise is the prevention of and intervention in cases of homelessness. Although unable to recall the precise numbers by heart, she firmly believes that Phoenix has a higher rate of homelessness than Boston, and says so. In principle, the accuracy of her judgments can be quantified in terms of their validity (i.e., the number of correct predictions divided by the total number of predictions she makes). Admittedly, listeners

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

can hardly infer an expert's judgment accuracy in this way given a single statement. They may, however, consult a different reference class by, for instance, calling up their impression of the average accuracy of experts they have seen in previous television debates.

The advantage of cues such as recognition, fluency, and likeability is that they are cognitively inexpensive. Pachur and Hertwig (2006), for instance, argued, and reported evidence, that the retrieval of recognition information precedes the retrieval of a probabilistic object cue and poses little to no cognitive cost. Physical attractiveness, a key determinant of likeability (think of Kennedy vis-à-vis Nixon), can be assessed from a face in as little as 13 milliseconds (Olson & Marshuetz, 2005; chapter 16). Given that easily accessible cues such as recognition can be highly predictive (Goldstein & Gigerenzer, 2002; Hertwig, Herzog, et al., 2008), listeners may be justified—regardless of their motivation and processing capacity—in relying on these cues.

Prediction 3. Impact of communication modality: A verbal message is more likely to be processed lexicographically than is a written message.

This prediction is derived as follows: In the study of heuristics, Gigerenzer and Goldstein (1996; see also Hertwig, Barron, Weber, & Erev, 2004; Hertwig & Erev, 2009) proposed a distinction between “inference from givens” and “inference from memory.” Inference from givens encompasses situations in which all the relevant information is fully displayed to participants. For instance, in many classic probabilistic reasoning tasks, such as the “Linda” problem, the experimenter provides all the relevant information, and individual knowledge about, say, feminist bank tellers is considered to be irrelevant (see Hertwig, Benz, & Krauss, 2008; Hertwig & Gigerenzer, 1999). Inference from memory, in contrast, entails memory search either within the individual's declarative knowledge base or in the external environment, so constraints such as time pressure and information cost can be assumed to shape these inferences. Indeed, as Bröder (p.125) (2012) demonstrated

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

experimentally, memory-based decisions differ from those based on givens: Given the naturally occurring information costs in inference from memory (compared with inference from givens), people are more likely to adhere to a frugal lexicographic strategy like take-the-best.

Based on these findings, PPT predicts that the processing of a message will be a function of the communication modality. If the message is displayed to listeners in written form, and they may repeatedly peruse it at their own pace, compensatory strategies such as tallying (see above) or Franklin's rule (Gigerenzer & Goldstein, 1999) are more likely to be used than noncompensatory strategies. In contrast, when a verbal message is processed online, with no opportunity (apart from memory retrieval) to return to the message, arguments are more likely to be processed using noncompensatory strategies such as take-the-best.

Prediction 4. Going versus not going beyond the information given: Listeners who lack object cues (apart from those included in the messages) and listener cues (e.g., recognition and fluency) cannot go beyond the information given in the message. According to the notion of vicarious functioning, however, they can recruit speaker cues such as expertise, credibility, and likeability.

This prediction depicts a situation in which the listener cannot go beyond the information given in the message and so cannot resolve conflicts between speakers—as they occurred in the fictitious debate—by relying on her own knowledge. Yet such a listener can turn to the speaker's expertise, credibility, and likeability as cues for deciding whom and what to believe. Therefore, speaker cues are more likely to be used by listeners who are confined to the information given than by listeners who can enrich the given mental model with additional knowledge of their own.

Prediction 5. Effect of “matching lenses”: The larger the match between the speaker's message and the listener's lens model, ceteris paribus, the smaller will be the size of

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

the change in the listener's judgment as well as the likelihood of a change in judgment direction.

Neo-Brunswikian research on interpersonal conflicts (see Brehmer, 1976; Dhimi & Olsson, 2008) suggests that communication effectiveness may depend on the fit between the lens models of various users or group members. Applying this idea to our context, deliberative and judicial speech involves a listener who evaluates a message with the goal of arriving at a correct judgment. One indicator of a message's (p.126) effectiveness could then be the amount and direction of any resulting changes in judgment. Assuming this indicator, the effectiveness of the messages should therefore decrease as the match between the speaker's and listener's mental models increases. The reason is that the larger the match between the speaker's mental model (as revealed in her message) and the listener's preexisting mental model (i.e., cue knowledge), *ceteris paribus*, the more likely they will arrive at the same judgment (see chapter 10).

Probabilistic Persuasion Theory: Test of Predictions

In this section, we report two experiments that are intended to provide preliminary evidence for PPT by testing Predictions 1 and 2, as well as to illustrate how Predictions 3 through 5 could be tested in future investigations. The first experiment was conducted expressly for the purposes of this chapter; the other was run by Reimer and Katsikopoulos in 2004.

Do Speakers Prefer a Few Good Cues (Prediction 1)?

According to Prediction 1, rather than presenting all available arguments, speakers select arguments (cues) as a function of some measure of goodness. Our first experimental test of this prediction focused on preferences (Box 4-1) because preferences (rather than inferences) have been the “home turf” of persuasion research. We placed participants in the role of salespeople (i.e., speakers) whose task it is to recommend their product to a customer. In preference tasks, such as the choice between hypothetical job candidates or between mobile phones, it is difficult to measure the accuracy of cues in terms of their ability to predict real-world outcomes

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

(i.e., cue validity). Preferences, however, imply alternative measures of goodness, one of them being the *distinctiveness* of a cue.

“Distinctiveness” refers to the extent to which an object has a cue that sets it apart from other objects. The challenge for the speaker in a persuasion context is to select cues that help persuade the listener to endorse the object, product, opinion, or course of action that the speaker has predetermined. Distinctive cues are assumed to be persuasive. For illustration, consider five mobile phones and their values on the six attributes (in the preferential domain we speak of *attributes* rather than *cues*) displayed in Table 4-3 and used in our study (Box 4-1). The distinctiveness d of a positive attribute i can be defined as follows: the number of objects (phones) with a negative value on this attribute divided by the total number of objects minus 1 (i.e., $d_i = N^-(N-1)$). The distinctiveness of an attribute is thus defined as the proportion of objects that differ from a focal object in their (p.127) attribute values. The higher the score, the more distinctive the object is on the attribute. Table 4-3 shows the distinctiveness rates for each of the attributes relative to the first mobile phone (the Nvite 400). Distinctiveness is defined only for attributes that provide an argument in favor of a specific product (i.e., no distinctiveness measure can thus be calculated for attribute 6).

Box 4-1: Experimental Task

Cast in the role of a salesperson, participants ($N = 58$ undergraduate students) were asked, in two separate tasks, to persuade a customer to purchase a mobile phone and a camera, respectively. Participants read: “You are a salesperson in a store. A customer asks you for a recommendation of a mobile phone (camera). The customer has already preselected five products that do not exceed his budget.” Each of the five mobile phones (cameras) was described on six positive attributes; the attributes could be present (e.g., phone has an MP3 player) or absent. Participants were then asked the

following two questions: “Which of the five mobile phones (cameras) would you recommend to the customer?” and “How would you argue in favor of your recommendation?”

The order of the two tasks was counterbalanced, and each participant was randomly assigned to one of the two orders. The underlying information structure (product × attributes), displayed in Table 4-3, was the same for all participants and both tasks. Various grids were constructed by randomizing the order of the columns and rows as well as the assignment of each of the six attributes (such as “MP3 player”) to the information structure. As a consequence, the (subjective) importance of the various attributes and the fictitious product labels were counterbalanced. Consequently, every attribute (e.g., “e-mail”) could be the most, second, third, fourth, fifth, or sixth most distinctive attribute.

The information structure in the experiment was designed so that many participants would recommend the first mobile phone in Table 4-3 (Nvite 400). This phone has more positive attributes than any other phone. It also offers every single attribute that any other mobile phone has, as well as one that no other phone has (the distinctive attribute shown in Table 4-3 is the MP3 player; in the experiment, the distinctive attribute was randomized across all attributes). Indeed, the Nvite 400 was the single most frequently recommended product; in their role as a salesperson, nearly all participants (54 out of 58; 93%) endorsed it (and the corresponding camera). These endorsers could have tried persuading potential customers to choose this phone by mustering up to five arguments in its favor (i.e., the five positive attributes present), as 20 (37%) of them in fact consistently did for both products. But consistent with Prediction 1, most endorsers of the Nvite 400 (34; 63%) indicated that they would not use all available arguments in attempting to persuade customers to buy it; on average, they selected only 3.8 out of the 6 attributes (across the two tasks 11, 10, 19, 20, 45, and 3 participants selected 1, 2, 3, 4, 5, and 6 attributes, (p.128)

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

Table 4-3: A Choice Set of Five Mobile Phones, Described on Six Positive Attributes, and the Attributes' Distinctiveness (Y = Attribute Present; N = Attribute Not Present)

Phone Attribute	Nvite 400	B-smart 24	GM Atlantic	Andersen 500	Wonee A20	Distinctiveness
1. Email	Y	N	N	N	N	1
2. Camera	Y	Y	N	N	N	0.75
3. MP3 player	Y	Y	Y	N	N	0.50
4. Talk time (9h)	Y	Y	Y	Y	N	0.25
5. Wireless (WAP) technology	Y	Y	Y	Y	Y	0
6. Video conferencing	N	N	N	N	N	—

Note. The distinctiveness rates were computed relative to the Nvite 400.

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

respectively). That is, even in a persuasion context in which speakers could select all attributes at no cost (they were given and did not need to be retrieved), most of them winnowed down the set of arguments.

Now we can ask: Among those participants who winnowed down the set of attributes, was the selection guided by the attributes' distinctiveness? Figure 4-4 shows the proportion of the time that participants selected each of the six attributes averaged across both tasks. The most distinctive attribute (Table 4-3) was selected more than 90% of the time; the least distinctive attribute was selected only about 65% of the time; and attribute 6, which described a feature that none of the objects possessed, was almost never mentioned.

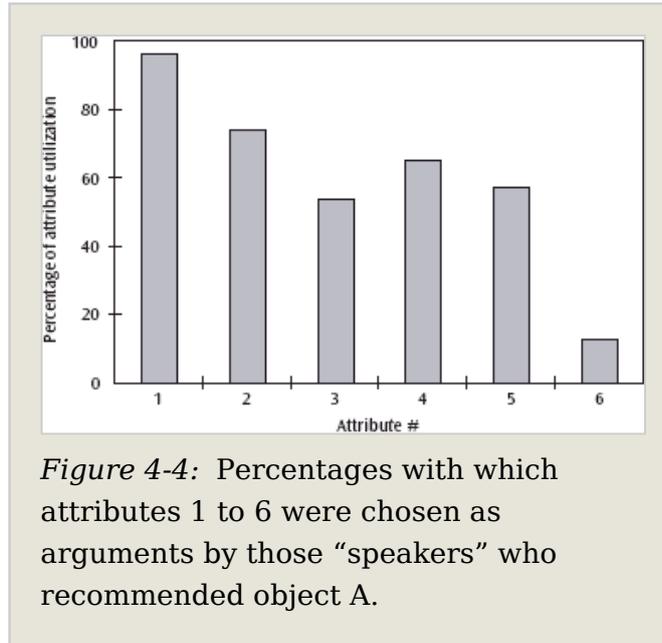
In sum, we found that, consistent with Prediction 1, most speakers do not present all available arguments in promoting a specific object, and that their selection of arguments mirrors one plausible measure of goodness: namely, the distinctiveness of attributes. Which measure of goodness a speaker will focus on depends on, among other variables, the context in which communication takes place. In related work, we could demonstrate that speakers also obey the relevance principle in selecting arguments. When the context rendered some attributes more relevant than others, participants chose the more relevant attributes more often than they did in a situation that did not provide hints about relevance.

The Primacy of Cognitively Inexpensive Cues (Prediction 2)

To test Prediction 2, we turn to an experiment conducted by Reimer and Katsikopoulos (2004). Consider a three-member search committee that must decide which of two final job candidates to invite for an interview. Using the available information, each member alone (p.129)

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

selects a favored candidate. Then the committee enters a negotiation stage in which all three members attempt to persuade the others of their respective choices. Structurally, this was the situation in which the



German participants in Reimer and Katsikopoulos's experiment found themselves. Specifically, their task was to find the correct answer, first individually and then in the three-person group, to questions such as “Which of these two U.S. cities has more residents: San Diego or San Antonio?” One straightforward strategy for reaching group agreement would be to settle on the opinion of the majority of group members (the majority rule; see Gigone & Hastie, 1997a).

Now consider the following conflict that sometimes arose at the negotiation stage. Two group members had heard of both cities and concluded independently that city *A* has more residents. The third group member, however, had heard of *B* but not *A* and concluded that *B* is larger on the basis of the recognition heuristic, which for such tasks is simply stated (Goldstein & Gigerenzer, 2002):

If one of two objects is recognized and the other is not, then infer that the recognized object has the higher value.

After the three members concluded their negotiations, what opinion prevailed? Given that two of the three group members deemed *A* to be larger and apparently could muster some knowledge about the two cities—that is, object cues such as those shown in Table 4-2—one might expect them to be able easily to persuade the third person that the correct answer is city *A*. In other

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

words, given that this person (p.130) has never even heard of one of the two cities under consideration, one might think that general knowledge about the cities would be more persuasive than an argument based merely on recognition or lack thereof. This is not what happened in Reimer and Katsikopoulos's (2004) experiment. In more than half of all cases (59%) in which two people recognized both cities and one person recognized only one, the opinion of the least knowledgeable person prevailed. That is, the least knowledgeable person succeeded in persuading one or both of the more knowledgeable group members to go along with his opinion. How can that be? All three members were equally motivated and had normal cognitive capacity. How, then, could knowledge about object cues be trumped by (partial) ignorance of the objects themselves? Within PPT, the explanation is simple: argument quality. Before creating the three-person groups, Reimer and Katsikopoulos quizzed respondents individually to find out which of 40 U.S. cities they recognized. The responses allowed the authors to estimate the recognition validity α (i.e., the cue validity for recognition knowledge) for each individual by calculating the proportion of correct inferences he would make if he used the recognition heuristic in all those pairs of cities where he had heard of only one city. Participants were then asked to perform the population comparison task for the pairs where they recognized both cities. From the answers, Reimer and Katsikopoulos could estimate a person's general knowledge validity β to be the proportion of correct responses for these pairs. The averages of the individual parameter estimates were $\alpha = 0.72$ and $\beta = 0.65$. In other words, in this environment, people who could employ the recognition heuristic simply had not only the better but also, at the same time, the more persuasive argument than those who relied on their general knowledge. And, indeed, the groups that went along with the opinion of members who could employ the recognition heuristic fared better than groups that adopted the opinion of the most knowledgeable members, who recognized both cities and therefore could not use the recognition heuristic (75% vs. 62% correct inferences).

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

The Impact of Communication Modality (Prediction 3)

Although we have not yet tested Prediction 3 experimentally, past research speaks to the presumed rationality or irrationality of heuristic and cue-based processing in different communication modalities (e.g., Chaiken & Eagly, 1976; 1983). (This research does not address the compensatory vs. noncompensatory nature of the processing.) Chaiken and Eagly (1983), for instance, concluded that:

videotaped and audiotaped modalities enhance the salience of communicator-related information [likeability] with the consequence that (p.131) communicator-related information exerts a disproportionate impact on persuasion when messages are transmitted in videotaped or audiotaped (vs. written) form. (Chaiken & Eagly, 1983, p. 241)

Two points are interesting here. First, the example of the Nixon-Kennedy debate seems to challenge Chaiken and Eagly's collapsing of the videotaped and audiotaped modalities. The debate's differential impact on the television audience (videotaped) and radio audience (audiotaped), respectively, suggests that the visual channel conveys presumably nonverbal cues that the auditory channel does not. That is, the audiovisual modality offers the largest repertoire of cues, followed by the acoustic channel, and lastly the written medium.

Second, according to the authors, reliance on communicator cues suggests heuristic processing, "whereas the relatively greater salience of message content in the written modality favors systematic processing" (Chaiken & Eagly, 1983, p. 254). Along with this inference comes their evaluative conclusion; namely, that the impact of communicator cues is "disproportionate" relative to that of what they describe as "message-based cognition." A decade later, Ambady and Rosenthal (1993) started a fascinating line of research that has by now established the power of "thin slices," or samples, of social behavior. Specifically, they found that undergraduate participants could predict college teachers' overall end-of-semester evaluations (an ecologically valid criterion variable, according to Ambady & Rosenthal, 1993, p. 431) on the basis of thin slices of the teachers' nonverbal behavior during instruction (i.e., silent video clips of less than 30, 15, and 6

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

seconds, respectively). In other words, communicator (as well as other) cues, processed in the blink of an eye, can be highly predictive of important target variables—a finding that contradicts dual-process theories' characterization of heuristic and cue-based processing as second-rate operations.

A Brunswikian Perspective and Probabilistic Persuasion

The framework proposed here draws heavily on the Brunswikian notions of the lens model, proximal cues, and vicarious functioning. The lens model allowed us to define an objective measure of cue goodness and thus of argument quality. Moreover, we proposed to model the cognitive processes that make up the Brunswikian lens in terms of fast and frugal heuristics (see Gigerenzer & Kurz, 2001) that process cues and arguments in a noncompensatory fashion, particularly when the communication mode is verbal.

According to Brunswik (1956), any organism has to cope with an environment full of uncertainties. Uncertainty certainly reigns in the world of human communication and social influence through (p.132) persuasion. For instance, the bulk of familial, public, and professional debates—Should you spend your vacation in location A or B? Should you go to college A or B? Should you hire candidate A or B?—require implicit or explicit predictions about the future. These predictions about the future are uncertain and require probabilistic inference strategies. Moreover, because of these uncertainties, the phenomenon of persuasion is probabilistic in nature and demands a probabilistic framework.

By bringing the Brunswikian lens model together with simple inferential heuristics, we can draw from a rich repertoire of inductive strategies for processing arguments, depending on the cognitive task at hand (e.g., choice, estimation, classification, and preference; see Gigerenzer et al., 2011). In analogy to the decoding of linguistic meaning (Wilson & Sperber, 2004), we conjecture that the listener's default processing of arguments is lexicographic and cue-based. Cues and therefore arguments come in different shapes and sizes. PPT declines to treat some categories of cues—specifically,

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

what are referred to in dual-process models of persuasion as “peripheral cues”—as second-class information and heuristic processing as suboptimal. Cognitively inexpensive information—such as recognition, fluency, and the speaker's perceived expertise—can be as or even more predictive than cue knowledge that is effortfully retrieved, and heuristic processing can be as or even more accurate than complex statistical procedures at making inferences (Gigerenzer & Brighton, 2009).

Many aspects of PPT need further expansion, development, and testing. Open issues include the following:

1. Which reference class is activated in a listener's mind? Argument validity can change substantially with the reference class the speaker and the listener have in mind (Hoffrage & Hertwig, 2006).
2. Are arguments always selected according to some measure of goodness? Alternative models of argument selection could consider other selection criteria, such as how widely shared, known, and reiterated an argument is (chapter 11; see also Hertwig, Gigerenzer, & Hoffrage, 1997).
3. Under what conditions do listeners compare the cues employed by the speaker against the ones stored in their memory and even enrich them by retrieving novel ones? For instance, do cues such as credibility and expertise (or lack thereof) trigger such verification processes?
4. If people rely on fast and frugal heuristics to process arguments—replacing cue integration with cue substitution—how well adapted will their final judgments be?

(p.133) Last but not least, the Brunswikian framework has important methodological consequences for persuasion research. The notion of vicarious functioning is closely related to Brunswik's great methodological innovation, which he called *representative design* (Brunswik, 1955; Dhimi, Hertwig, & Hoffrage, 2004). In *systematic design*, experimenters select and isolate one or more independent variable(s) by varying

Probabilistic Persuasion: A Brunswikian Theory of Argumentation

them systematically while either holding extraneous variables constant or allowing them to vary randomly. Brunswik opposed this experimental approach on the grounds that it risks destroying the natural causal texture of the environment an organism has adapted to (Brunswik, 1944) and “leaves no room for vicarious functioning” (Brunswik, 1952, p. 685). In other words, Brunswik argued that systematic design obliterates the very phenomenon under investigation or at least alters the processes underlying it in such a way that the results obtained are not representative of people's actual functioning in their ecology. In representative design, experimental stimuli are representative of a defined population of stimuli with respect to the numbers, values, distributions, intercorrelations, and ecological validities of their variable components (Brunswik, 1956).

The debate about systematic versus representative design (see Dhimi et al., 2004) should not be mistaken for an obscure academic quarrel. As in other areas of experimental psychology, systematic design (and its sophisticated variants such as factorial design) remains the preferred method of research on persuasion, in which cues ranging from expertise, credibility, and likeability, to object cues, are artificially decoupled and systematically orthogonalized, thus unraveling the probabilistic texture of the environment (for a similar critique of contemporary research in social perception, see Funder, 2001). In our view, entrenched dichotomies such as peripheral cues versus arguments and heuristic versus systematic processing are entangled with the use of systematic design, which presumes and fosters theorizing in terms of dichotomies. Experimental environments must retain the environment's probabilistic texture in order to shed light on people's actual functioning in it. If we dare to complement the use of systematic design with representative design, counterintuitive and surprising discoveries await us—such as the discovery that lack of recognition sometimes has higher validity than people's general knowledge, and that partial ignorance thus can be justifiably more persuasive than knowledge (Reimer & Katsikopoulos, 2004).

Conclusion

The dual-process approach to persuasion has undeniable merits. It has identified a number of systematic relationships between diverse determinants of persuasion. It can also accommodate a wide range of (p.134) empirical findings in a simple theoretical framework. Notwithstanding their explicit framing as accounts of distinct information-processing modes, however, dual-process models fail to specify the processing of what they see as distinct: namely, arguments and cues.

Thomas Kuhn (1962) emphasized that comparisons of established and new theories are complicated by (among other things) the incommensurability of parts of their lexica, and that such difficulties open up the discourse on theory choice to the influence of persuasion (Kuhn, pp. 93, 152). This view has often been read as implying a lack of any good reasons for choosing a new theory over an old one. In later work, however, Kuhn (1970) clarified his meaning: "To name persuasion as the scientist's recourse is not to suggest that there are not many good reasons for choosing one theory rather than another" (p. 261). We hope that in this chapter we have succeeded in providing some good arguments to adopt a Brunswikian framework for modeling the cognitive processes that underlie persuasion.

Notes:

(¹) Throughout the chapter, we use the term *speaker* rather than *communicator*. With this choice of words, we focus on verbal communication (but see Prediction 3).

(²) Hammond (1955; Hammond, Stewart, Brehmer, & Steinmann, 1975) and Brehmer (1976) extended and adapted the lens model to the study of social judgment, interpersonal conflict, and group decision making (for a collection, see Hammond & Stewart, 2001; Gigone & Hastie, 1997a). Using the lens model, Burgoon, Birk, and Pfau (1990) analyzed the relationship between nonverbal cues and a speaker's persuasiveness and credibility.

Probabilistic Persuasion: A Brunswikian Theory of Argumentation



Access brought to you by: Max-Planck Society