



Series Editor: **Derek J. Koehler**, *University of Waterloo, Canada*

The purpose of the *Society for Judgment and Decision Making* book series is to convey the general principles and findings of research in judgment and decision making to the many academic and professional fields to which it applies. Because the field of judgment and decision making is largely a formal one (similar to mathematics), its principles and findings are applicable to a wide range of disciplines, including psychology, medicine, social policy, law, management science, economics, and accounting.

The books in this series are aimed at researchers and their upper-level students. Most of the books are multiauthored volumes written by authorities in the field and sponsored by the Publications Committee of the Society for Judgment and Decision Making.

PUBLISHED

Social Psychology and Economics, De Cremer, Zeelenberg & Murnighan
The Science of Giving: Experimental Approaches to the Study of Charity,
Oppenheimer & Olivola
Perspectives on Framing, Keren
A Handbook of Process Tracing Methods for Decision Research,
Schulke-Mecklenbeck, Kuehberger & Ranyard

For continually updated information about published and forthcoming titles in the *Society for Judgment and Decision Making* series, please visit www.psypress.com/sjdm

Perspectives on Framing

Edited by
Gideon Keren
University of Tilburg, The Netherlands

2011

2

Levels of Information: A Framing Hierarchy

SHLOMI SHER and CRAIG R. M. MCKENZIE

INTRODUCTION

Framing experiments seek to rigorously separate out the effects of relevant and irrelevant information on human judgment and choice processes. Because they appear to elegantly streamline the normative analysis of human cognition, these experiments have assumed a central place in the so-called “Rationality Debate” – the controversy, within and between the various social sciences, over the rationality of human action (Shafr & LeBoeuf, 2002). As Kahneman (2000b, p. xv) has argued, framing effects “provide a compelling reason to separate descriptive from normative models of choice. It is surely rational to treat identical problems identically, but often people do not.”

The objective of this chapter is to characterize the power and limitations of framing as an experimental tool, in relation to the Rationality Debate. The normative analysis of framing effects, we argue, is more complex than is often supposed – it ultimately depends on a formal concept of “information” and a view of human cognitive systems, in relation to one another and to the information environments in which they usually operate. In the pages that follow, this argument is developed in the context of a framework for thinking and talking about framing research – its vocabulary, its goals, and its normative interpretation. Accordingly, we begin with a brief overview of the concepts of framing, the uses of frames, and the analysis of framing effects. We then outline the general plan of the chapter.

Concepts of Framing

“Framing” is a widely and sometimes loosely used concept that refers to situations in which a speaker, often with a persuasive agenda, selects one among multiple possible ways of presenting “the same information” to a listener. For example, a retailer may describe a ground beef product as “25% fat” or “75% lean”. “Framing effect” refers to a class of well-established experimental phenomena in which

people are found to respond differently, on average, to different descriptions that convey "the same information" in different ways. For example, experimental subjects have been found to evaluate ground beef more favorably when it is described as "75% lean" (Levin, 1987; Levin & Gaeth, 1988). These experimental effects are generally thought to violate a normative invariance principle, which requires identical responses to equivalent descriptions (Tversky & Kahneman, 1986). Note that the above definitions of "framing" and "framing effect" refer to "information" – a concept that also needs a definition. In fact, to fully capture its use in normative analysis, it needs multiple definitions. The various meanings of "information" are detailed later in this chapter.

The term "frame" is used inconsistently in the literature. Druckman (2001; see also Chapter 14 of the present volume) collected several influential but incompatible definitions, and drew a conceptual distinction between "frames in communication" and "frames in thought". A frame in communication is simply the description delivered by a speaker to a listener. In a framing study, the frame in communication is the experimental manipulation – the specific wording conveyed to the subject in, for example, a questionnaire. A frame in thought, by contrast, is a psychological perspective on a situation, a way of looking at things. In a framing study, a frame in thought may be the theoretical mechanism whereby the frame in communication is proposed to influence the subject's response. It is important to underline the simple fact that research designs never manipulate frames in thought directly – if they do, they do so indirectly, by way of a frame in communication. We will use the term "frame" here exclusively to refer to the frame in communication – the overt wording received by a listener in a persuasive setting, or by a subject in a framing experiment.

Information Analysis and Psychological Analysis

The normative evaluation of framing effects may assume either of two related forms – an information analysis of frames, or a psychological analysis of processes. An information analysis seeks to establish that the frames in an experiment carry "the same information", and that they therefore fall under a normative invariance principle requiring identical responses to equivalent descriptions. When an information analysis succeeds in equating frames, different responses may underlie an important additional assumption noted below) be classed as intrinsically incoherent, even if the psychological processes producing the responses are not understood. When an information analysis fails to rigorously equate frames, and we will argue that this happens more often than is often thought, it is still possible to ask normative questions about the psychological processes that mediate their effects. The psychological analysis of an effect asks whether and how the cognitive processes underlying it make use of relevant information in appropriate ways – and why they don't when they don't.

The first part of this chapter develops an information analysis of framing effects. What does it mean to say that two frames carry "the same information", and when and why does this matter? First, we consider the logic of "the equivalence method" – the line of reasoning by which framing researchers reason from

their experimental effects to conclusions about human irrationality. We then introduce a general concept of "information", and delineate five "levels of information" relevant to the analysis of decision and belief. The information analysis of frames, we argue, turns out to be a problem of considerable empirical complexity, requiring a study of the natural communication environments in which frames are typically selected by speakers. Against the background of this information analysis, the chapter then turns to the psychological analysis of framing effects. We outline two psychological perspectives on framing, in terms of a tentative "two-systems" picture of human cognition (Evans, 2003; Kahneman & Frederick, 2002; Stanovich & West, 2000). These two psychological perspectives, in turn, suggest conflicting conceptions of the rationality of intuitive judgment in complex information environments.

THE EQUIVALENCE METHOD

Experimental studies of framing fall within a tradition of research that bridges two different kinds of theory – normative theories that specify the conditions of rational choice and inference, and empirical theories that specify the conditions of actual choice and inference. This research tradition – the psychology of judgment and decision making (Gilovich, Griffin, & Kahneman, 2002; Kahneman, Slovic, & Tversky, 1982; Kahneman & Tversky, 2000) – seeks, in part, to determine whether the empirical facts of human psychology line up with the normative requirements of logic, probability, and decision theory. Failures of alignment have potentially broad implications for social scientists who apply normative theories in modeling individual and collective choice phenomena (e.g., Shwartz, 2004), as well as for our general understanding of the foibles of friends and the madness of crowds (e.g., Gilovich, 1991). Despite its obvious interest and importance, however, this tradition encounters a formidable obstacle from the outset.

Nearly all interesting human judgment and choice situations are far too complicated for the explicit calculation of optimal responses from normative theory. Therefore, the simple research strategy of comparing observed responses with computed optima is generally unavailable to the researcher – the optima cannot be rigorously derived.

Researchers in judgment and decision making have developed an arsenal of clever strategies to circumvent this problem of normative hypercomplexity. Many of these strategies are coherence based (Kahneman, 2000a). These strategies exploit the fact that the mutual incoherence of a collection of imperfect responses is often easier to establish than the non-optimality of any response in the collection. The equivalence method – in which the presence of irrational responses somewhere in a response set is signaled by the presence of a framing effect in the aggregate – is a leading coherence-based experimental strategy.

The equivalence method has three ingredients. *Two different descriptions*, which according to accepted normative theory *convey the same information*, are presented to two groups of subjects. The researcher then determines whether the two groups form *the same judgments and decisions on average* in response to the

different descriptions. When they don't, the researcher concludes that at least some people in the sample are judging and deciding in a non-optimal manner.

This conclusion may seem unassailable, but it is useful to examine its basis explicitly. Consider the usual situation, in which the researcher is unable to formally compute the optimal response to any description used in the experiment. Indeed, the researcher will generally be willing to allow that different subjects may have different optimal responses, as subjects will vary uncontrollably but acceptably in both background knowledge and values. Nevertheless, for a fixed subject in the experimental population, the optimal responses to the two possible descriptions must be identical, because the information the descriptions convey is the same. Therefore, in any large sample drawn randomly from a population of optimal responders, the obtained distribution of responses (whatever it happens to be) should be identical across the two description conditions. If the two response distributions differ meaningfully, then there must be non-optimal responders in the sample, even though we are unable to point to any particular non-optimal response.

The above line of reasoning makes two important but non-trivial assumptions. First (*the uniqueness assumption*), it presupposes that for each subject there is a unique optimal response to each description. That is, if one response in the response set is at least as good as all other responses, then it is strictly better than all other responses. Without this assumption, we could not exclude the possibility that different descriptions systematically bias some subjects' arbitrary choices between normatively indistinguishable responses, leading to innocuously shifted response distributions. Second (*the equivalence assumption*), the above line of reasoning presupposes that the unique optimal response to any description must be the same as the unique optimal response to any other description that conveys "the same information".

The uniqueness assumption is rarely made explicit in experimental treatments of human rationality, though it often plays an important hidden role. For the purpose of this chapter, we will generally assume that the uniqueness assumption is satisfied – that is, that for each subject there is one and only one optimal response to each description – not because we believe this assumption to be generally warranted (see Shafer, 1988; and the contributions in Chang, 1997), but because we want to isolate the equivalence assumption for analysis.

As we show in the following pages, the normative viability of the equivalence assumption – the assumption that the unique optimal responses to two framings of "the same information" must be identical – depends on the level of information at which one is operating.

INFORMATION

Concepts of information are indispensable everywhere, from thermodynamics to political science. Although they share a common core, the information concepts used by the psychologist, the logician, the economist, and the sociologist differ in significant ways. This conceptual variation across disciplines rarely creates problems within disciplines. But this variation complicates the connections we would like to

draw between different theories (e.g., deductive logic and experimental psychology) that operate at different levels of information. This section presents a simple framework that captures what the various information concepts have in common, and specifies three parameters on which they characteristically vary.

Information involves a relation between two sets – (1) a set of "signals" and (2) a set of possible "states of the world". Colloquially, the signals are the "carriers" of information: they commonly are symbol sequences in some compositional language, but they don't need to be. Formally, a signal carries information by virtue of (3) a presumed underlying mapping, "the information function", which assigns to each signal a unique distribution over the possible states of the world.

Various information concepts capture the many ways in which receivers can use signals to learn about the world. Thus an information function may be *probabilistic*, mapping each signal to a distribution assigning real numbers between 0 and 1 to every possible state, or *deterministic*, mapping each signal to a distribution assigning 1 to some states (the states consistent with the signal) and 0 to all others (the states not consistent with the signal). Which information concept is most useful in analyzing a given domain depends on the structure of the domain and the aims of the analysis. For example, information concepts will be specified differently by the author of an English dictionary, who seeks to prescribe and codify normal usage, and by the vision scientist, who seeks to understand the information potentially contained in arrays of retinal photoreceptors. Their information concepts will implicitly presuppose different signals (English words *versus* activated retinal arrays), different world states (ideas and situations *versus* surface configurations), and different information functions (a deterministic function that captures canonical meaning *versus* a probabilistic function that captures possible inference).

Information concepts are useful wherever signals stand for states, but the range of such situations is richly varied and the range of information concepts is correspondingly diverse. A particular information concept is formed when three parameters are specified: the set of signals, the set of states, and the function that maps the former to distributions over the latter. These three parameters are routinely set differently in different domains and at different levels of analysis, to suit the varying purposes of the analyst. The next section considers the information concepts that are potentially relevant to the analysis of frames and their experimental effects.

LEVELS OF INFORMATION

In this section, five information concepts are introduced. Each concept is well suited to a particular level of analysis, and each supplies a sense in which a pair of signals can be regarded as equivalent – that is, as carrying "the same information". These five levels of information are summarized in Table 2.1. Each level uniquely specifies the three parameters described above: the signals that are permitted, the states of the world that are considered, and the mapping from signals to states that is contemplated. The first level of information (*unrestricted inference*) has the

TABLE 2.1 Levels of information.

Level of analysis	Information function	Standard of equivalence	Definition	Appropriate for
1. <i>Unrestricted inference</i>	Signal → Probability distribution over world states conditional on signal selection	Information equivalence	Nothing can be inferred from the fact that utterance A, rather than utterance B, was encountered in this environment	The analysis of psychology experiments, involving utterances A and B
2. <i>Deductive analysis</i>	Proposition → Set of world states in which the proposition is true	Logical equivalence	Nothing can be inferred from the fact that utterance A is true that can't be inferred from the fact that utterance B is true (and vice versa)	The formal analysis of deductive arguments, involving premises A and B
3. <i>Economic analysis</i>	Situational description → Idealized agent belief-value state (risk profile)	Formal equivalence	In two idealized agent belief-value states A and B, all the probabilities of all of the outcomes in A and B are identical	The idealized analysis of economic situations, involving risk profiles A and B
4. <i>Data analysis</i>	Data summary → Collection of data sets that the summary accurately summarizes	Observational equivalence	The empirical data presented in format A, if accurate, could also be presented in format B (and vice versa)	The analysis of arbitrarily formatted empirical data, depicted in form A or B
5. <i>Content analysis</i>	Messy, variable across domains Argument highlighting facts, assessments → Unweighted set of facts, assessments	Substantive equivalence	Communications A and B make the same contribution to the listener's "relevant" store of facts and evaluations	The analysis of arguments A and B in particular domains (e.g. political persuasion)

least structure and the greatest generality: at this level, the information content of a signal encompasses everything that can be legitimately inferred from the fact that it was received. This, we argue, is the level at which the rationality of real human responses in experimental and natural environments should ultimately be evaluated. The four lower levels of information are more tightly structured, owing to specialized restrictions on the set of world-states and/or on the mapping from signals to states. Each of these four information levels corresponds to a traditional level of normative analysis, and each has been used to define equivalence of information in experiments employing the equivalence method.

Level 1: *Unrestricted Inference*

Consider the predicament of an ideal consumer of information receiving signals in a conversation, a marketplace, or a psychology experiment. From the perspective of the consumer, these signals may be conceived as evidence about the state of an uncertain world. At Level 1, two signals are "information equivalent" if and only if they supply identical evidence – that is, if the receipt of either signal would lead the ideal consumer to update her beliefs in the same way. The formal parameters of this information concept are set as follows.

The states of the world are construed very broadly, to include the richest available catalogue of possible empirical situations; and the signals are viewed as the output of a particular signal-selection process operating in the real world. The recipient of the signal is assumed to have some partial prior beliefs about the state of the world and about the signal-selection process. For convenience, we assume that the recipient's uncertain model of the world and the selection process can be represented probabilistically – that is, by a probability distribution over states of the world, and conditional probabilities for the production of different signals in different possible states of the world. Then the recipient's probabilistic model of the forward mapping from world-states to selected signals permits the construction (e.g., by Bayes' theorem) of an inferential backward mapping from received signals to probability distributions over the world-states in which they might have been generated. The "unrestricted information content" of a signal is defined as the optimally updated probability distribution over world-states, conditional on receipt of the signal from the signal-selection process.

In a natural framing environment, descriptions, the signals, are selected on the fly by a human speaker, the signal-selection process. The selection process is presumably non-random – human speakers do not toss coins in choosing to describe ground beef as "25% fat" or "5% lean". Listeners, in turn, bring a lifetime of conversational experience to the interpretation of the novel descriptions selected by speakers. When a listener updates her beliefs on the basis of a speaker's description, this updating process may reflect a rich store of implicit knowledge about the conditions under which speakers select different kinds of description.

Formally, two signals A and B are "information non-equivalent" if they supply relevantly different information at Level 1 – that is, if there is a state of the world C satisfying the following two conditions:

1. *Choice relevance*: A decision maker's beliefs about C can legitimately affect her decision.
2. *Differential inference*: The probability of C conditional on receipt of signal A (denoted " $P(C|A)$ ") differs from its probability conditional on receipt of signal B (" $P(C|B)$ ").

If such a condition C exists, some decision makers may legitimately make different choices upon receiving the different signals A and B. How can we empirically demonstrate that a particular choice-relevant background condition satisfies the differential inference condition (2)? It suffices to show that it satisfies:

3. *Differential production*: Among those cases in which the signal-selection process selects either A or B, the selection process is more likely to select A when C holds (" $P(A|C)$ ") than when C fails (" $P(A|\neg C)$ ").

Differential production (3) implies differential inference (2). To see this, consider a signal-selection process S that sometimes (i.e., with non-zero probability) selects signals A and B, and restrict attention exclusively to those instances in which one of these two signals is selected. We use $\neg A$ to denote the case in which the selection process does not select A – that is, in which it instead selects B. Let C denote a choice-relevant background condition that may or may not ($\neg C$) obtain. For example, S may be a speaker truthfully describing a particular medical treatment, with frame A = "The treatment leads to 75% survival after 5 years", frame B = "The treatment leads to 25% mortality in 5 years", and C denoting the state of the world in which the treatment in question leads to more deaths than the available alternatives.

Now suppose differential production: $P(A|\neg C) < P(A|C)$. That is, in those cases where either A or B is selected, A is more likely to be selected when C holds than when C fails. For example, we may have experimentally determined that speakers are more likely to select "mortality" descriptions for treatments leading to relatively many deaths (cf. Level 2 below). Because $P(A)$ is a weighted average of $P(A|\neg C)$ and $P(A|C)$, this implies that $P(A) < P(A|C)$. This is equivalent to $1 - P(A|C) < 1 - P(A)$ and therefore:

$$\frac{1 - P(A|C)}{1 - P(A)} < 1$$

By Bayes' theorem, together with the above inequality, we have:

$$P(C|\neg A) = \frac{P(\neg A|C) P(C)}{P(\neg A)} = \frac{1 - P(A|C)}{1 - P(A)} P(C) < P(C)$$

But, since $P(C)$ is a weighted average of $P(C|\neg A)$ and $P(C|A)$, this means that $P(C|\neg A) < P(C|A)$. That is, C is more likely to hold when A is selected than when B is selected – the differential inference condition is satisfied. Since C is choice

relevant, the speaker's choice of frame may influence the listener's decision. In our example, if speakers are more likely to select "mortality" descriptions for treatments with relatively high rather than low mortality, a hypothetical ideal information consumer would evaluate the treatment less favorably upon receipt of the "25% mortality" frame than she would upon receipt of the "75% survival" frame.

In the normative evaluation of experimental effects, we will generally assume that subjects interpret signals as they would in the closest natural signal-selection (i.e., communication) environment. Of course, the signal-selection process in experimental design is very different from those that operate in the "natural world"; but the stimulus-selection process in experiments is typically concealed from the subject's view, and in many experiments cooperativeness requires the subject to treat artificial stimuli roughly as they would in corresponding natural situations.

Equivalence method experiments are usually devised at Levels 2–5, detailed below. Consequently, frames traditionally classed as currying "the same information" may differ relevantly in their unrestricted information content. When this happens, the selection of a frame by the selection process is said to "leak" choice-relevant information (Sher & McKenzie, 2006). The ways in which frames equivalent at Levels 2–5 may leak information at Level 1 are illustrated in the discussions that follow:

Level 2: Deductive Analysis

Framing researchers rarely explicitly stipulate the level of information at which a pair of frames submitted to the equivalence method must be equivalent. However, it is occasionally said that framing effects are counter-normative when the frames are *logically equivalent* (e.g., Johnson-Laird & Shaflir, 1993; Rubinstein, 1998; Shaflir, 1993).

The logical information content of a statement plays a critical role in the normative analysis of arguments in which it figures. Deductive analysis is concerned with inference procedures from sentences to sentences that are guaranteed to preserve truth value. Accordingly, information concepts relevant to deductive analysis are characterized by weak constraints on the set of possible world-states (to allow generality) and strong constraints on the information function from signals to states (to ensure necessity). With minor simplification, the logical information concept can be characterized by the following parameter settings.

The set of world-states is construed in the widest possible way, to include formal and merely possible as well as actual empirical objects: logical principles equally constrain arguments about triangles, solar systems, and traffic regulations. Logical vocabularies permit the construction of sentences with truth value the signals. These vocabularies often include simple logical particles "and", "or", etc., which are given more rigid interpretations than the same words receive in normal conversational language. Finally, deductive analysis is characterized by an information function that, in essence, maps each sentence to the subset of possible states of the world in which the sentence would be true. This collection of states is the sentence's "logical information content". The states in this collection are not

probabilistically weighted, because deductive analysis covers every possible scenario. Two sentences are “logically equivalent” if they semantically entail each other – that is, if each is true in every possible state of the world in which the other is true.

A major division of the framing literature concerns logically equivalent descriptions of proportion. In *attribute framing* (Levin, Schneider, & Gaeth, 1998), the value that a single object assumes on a single bounded dimension can be truthfully described in terms of either of two logically interchangeable attribute proportions. Typically, these attributes are charged with opposite affective valence. For example, a medical treatment has “75% survival” across a given time window if and only if it has “25% mortality” within the same window (e.g., Levin, Schnitger, & Thee, 1988; Martean, 1989; Wilson, Kaplan, & Schneidermann, 1987); a particular parcel of ground beef is “75% lean” if and only if it is “25% fat” (Levin & Gaeth, 1988); and so on. The standard finding in attribute framing experiments is a “valence-consistent shift”: subjects evaluate objects less favorably when they are described in terms of the negative attribute (Levin et al., 1998). Medical treatments seem less attractive when they are described as leading to “25% mortality” rather than “75% survival”; “25% fat” beef seems less palatable than “75% lean beef; and so on.

Typical attribute frames are equivalent at Level 2 – in every possible world, state in which one attribute description is true, the other description must also be true. However, are these logically equivalent descriptions information equivalent at Level 1? Recall that two frames are non-equivalent at Level 1 if there is a choice-relevant background condition such that typical speakers, choosing among the two frames, are more likely to select one of the frames when the condition holds than when it fails. In fact, experimental work has identified two such background conditions – reference points and implicit recommendations – that have broad relevance in attribute framing tasks.

First, speakers are more likely to describe objects in terms of attributes that exceed a relevant reference point – the prior, expected, or standard level of the attribute. This frame-selection regularity has been documented in multiple experimental tasks involving both neutral and valenced descriptions (McKenzie & Nelson, 2003; Sher & McKenzie, 2006). For example, subjects were more likely to describe a medical treatment in terms of its “mortality” rate, rather than its “survival” rate, when the treatment led to more fatalities and fewer survivors than the norm (McKenzie & Nelson, 2003). An option’s relationship to its likely alternatives – unspecified in typical attribute framing tasks – is broadly relevant to evaluation. Second, the selected frame may leak information about the speaker’s attitude – that is, a kind of implicit recommendation. Speakers with more negative attitudes toward an object are more likely to select negatively valenced attribute frames in describing it (Sher & McKenzie, 2006). The speaker’s choice of attribute frame is therefore potentially relevant to the listener’s evaluation whenever an explicit recommendation from the speaker could reasonably sway some listeners. Because negatively framed objects are likely to suffer in comparison with their alternatives, or in the eyes of the speaker, standard attribute frames differ in their Level 1 information content. A valence-consistent shift is compatible with the information analysis of typical attribute frames.

Level 3: Economic Analysis

A second important division of the framing literature investigates choices among gambles. In *risky choice framing* (Levin et al., 1998; reviewed in Kühnberger, 1998), subjects receive different descriptions of risk-taking scenarios involving the same probabilities and outcomes. It is sometimes said that framing effects violate rational norms when the frames are “formally equivalent” (e.g., Frisch, 1993; Kühnberger, 1998). The formal system that classes different descriptions of identical probability–outcome profiles as equivalent is the so-called “rational actor model” of classical economic theory.

Economic analyses of choice normally stipulate a special restriction on the possible states of the world. The world consists of idealized choice situations; in each choice situation, the decision maker’s representation of each alternative is assumed to be captured by a sequence of probability–outcome pairs $(P_1, O_1; P_2, O_2; \dots)$, where each P_i is the probability assigned to outcome O_i . Furthermore, to every outcome there corresponds a definite numerical valuation $U(O_i)$, the utility; the decision maker implicitly attaches to the i th outcome. Each alternative in every possible episode of choice under uncertainty is assumed to be captured by one such probability–outcome representation. Therefore, under any acceptable notation for describing choice situations in the economic model, each description is mapped onto a single probability–outcome representation, and this representation comprises the description’s “formal information content”. Two descriptions are “formally equivalent” when, in this way, they single out the same probabilities and outcomes.

The best-known risky choice framing effect is Tversky and Kahneman’s (1981) widely replicated “Asian disease problem”, detailed in Chapter 1. As predicted by prospect theory (Kahneman & Tversky, 1979), subjects’ preferences exhibit marked risk aversion when public health programs are described in terms of probabilities that different proportions of a population will be “saved”, and marked risk seeking when the programs are described in terms of how many will “die”.

Tversky and Kahneman’s 1986, p. S258) two-stage problem is another important risky choice framing effect. Some subjects were told to assume themselves “richer by \$300 than you are today”, and were then asked to choose between two gambles involving potential gains – either “a sure gain of \$100” or “50% chance to gain \$200 and 50% chance to gain nothing”. Subjects in a second condition were told to assume themselves “richer by \$500 than you are today”, and were then asked to choose between two gambles involving potential losses – either “a sure loss of \$100” or “50% chance to lose nothing and 50% chance to lose \$200”. Note that the two problem descriptions are formally equivalent when outcomes are recoded as changes in total wealth. In both conditions, the sure option results in a certain gain of \$400 over the subject’s initial level of wealth, while the gamble results in a 50% chance of gaining \$500 and a 50% chance of gaining \$300 over the starting level. In line with prospect theoretic predictions, subjects tended to prefer the sure thing when choosing among incremental gains, but preferred the gamble when choosing among incremental losses.

Logically equivalent descriptions of probabilities and outcomes are always

formally equivalent, but formal equivalence does not imply logical equivalence. Accordingly, the formally equivalent frames in risky choice framing studies sometimes are equivalent at Level 2, but often are not. For example, the frames in the Asian disease problem are logically equivalent (though see Jou, Shanteau, & Harris, 1996), while the frames in the two-stage problem are logically incompatible. The truth of each two-stage frame entails the falsity of the other, because different event sequences are stipulated, but the descriptions are not distinguished by classical economic analysis, because the same outcomes are reached.

The logical relationship between formally equivalent frames is important, because the normative analysis of frames requires stronger assumptions at Level 3 than at Level 2. To see this, assume that Level 1 leakage can somehow be excluded from the information analysis of frames – for example, that frames are known to be selected at random (see the next section). Then a Level 2 invariance principle, requiring identical responses to logically equivalent frames, assumes the existence, for each decision maker, of a unique best response to the common collection of situations singled out by the two frames. Level 2 invariance, however, presumes nothing about the factors that matter in evaluating this collection of situations. A Level 3 invariance principle is much more demanding, requiring that all choice-relevant beliefs can be captured in probability–outcome form and that, within this representation, summated objective consequences are the only factors that matter for rational choice. These more exacting assumptions leave a Level 3 invariance principle more vulnerable to normative criticism. For example, the consequentialism presupposed in Level 3 choice models has been alleged to lead to deep problems in moral decision domains (e.g., Williams, 1973). In our discussion of the two-stage problem below, we note a further normative complication arising at Level 3, but not at Level 2.

Frames, however, are not selected at random in natural environments. This raises the question: Are risky choice frames, equivalent at Level 3, and sometimes at Level 2, also equivalent at Level 1? Because the probabilities and outcomes in a standard risky choice description comprise multiple independently moving parts, the analysis of frame selection is considerably more difficult in risky choice than in attribute framing. Furthermore, superficially similar risky choice framing problems may receive qualitatively different information analyses at Level 1 as well as at Level 2. With these caveats in mind, we consider possible Level 1 features of the risky choice framing problems noted above.

It is unclear whether the two framings of the Asian disease problem are information equivalent at Level 1. Shier and McKenzie (2005) found that speakers spontaneously preferring the sure thing tended to use the survival frame to describe it; and, when subjects were overtly assigned persuasive agendas, they tended to frame the full problem in ways that would lead listeners to adopt their position. Similarly, van Buiten and Keren (2009) documented a general compatibility between the descriptions chosen by speakers and the choices made by listeners in this problem. However, while these findings provide suggestive evidence for Level 1 information leakage in the Asian disease problem, no general regularity in frame selection has been documented that (1) holds across multiple natural frame-selection domains and (2) differentially predicts the joint framing of the

sure thing and the gamble as a function of choice-relevant background information. It is important to add that this effect is predicted by prospect theory, which makes no reference to regularities in frame selection. Therefore, the relevance of possible Level 1 information to the psychological analysis of the Asian disease problem (considered later in “Two views of two systems” section) can only be assessed in this context.

The Level 1 information analysis of the two-stage problem is unlikely to be illuminating, for two reasons. First, there are few comparably natural ways to describe each of the monetary sequences in this problem. Second, framing effects of this kind may be strongly expected in arbitrary information environments. In the two-stage problem, the formal equivalence of final outcomes hinges on the initial imagined \$300–\$500 adjustment to the subject's wealth level. But people generally don't know their total personal wealth with \$200 precision. Therefore, people will be largely insensitive to \$200 perturbations in wealth level in any reasonable model of human decision making. For this reason, some framing effects of this general kind will almost certainly arise, even if patterns of description selection can be ignored. The two-stage problem in this way illustrates a clear violation of the Level 3 invariance requirement presupposed in classical economic theory. An argument can be made, however, that rational choice should not require omniscience, and in particular should not assume fictional total knowledge about personal wealth. Therefore, whether the two-stage effect exposes the irrationality of human decision making, or simply illustrates the normative inadequacy of classical economic analysis, is open to dispute.

Risky choice framing differs in two important respects from the other levels of framing considered here. First, when the formally equivalent frames in a risky choice study are non-equivalent at Level 2, the normative force of the Level 3 invariance requirement depends on relatively strong assumptions, potentially complicating the normative analysis of effects. Second, these studies often draw especially strong motivation from detailed psychological theory, and therefore the psychological analysis of risky choice framing effects must be situated in an especially rich empirical and theoretical context.

Framing at Level 3 extends beyond traditional risky choice framing to encompass any pair of choice problem descriptions in which corresponding options have identical consequences, even as other features of the described situation vary. In “default effects”, for example, subjects decide whether to take action in overriding an explicitly designated default choice. When the consequences of taking action (failing to act in Problem 1 are identical to the consequences of failing to act (taking action) in Problem 2, decision makers often exhibit a bias toward inaction across problem forms. This may help to explain why rates of organ donation are much higher in countries with a donation default than in countries with a non-donation default (Johnson & Goldstein, 2003). The problem descriptions in default effects are equivalent at Level 3; however, they need not be equivalent at Level 1. For example, McKenzie, Liersch, and Finkelstein (2006) reported evidence that the selection of public policy defaults leaks potentially choice-relevant information about policy makers' recommended course of action.

Level 4: Data Analysis

The equivalence method extends beyond the traditional framing literature (see McKenzie, 2004, for a general discussion). Researchers have developed systematic methods to describe and infer statistical relationships on the basis of bodies of observational data. Normative principles formulated at this level of analysis have been employed in the experimental evaluation of lay inference, where systematically different responses to observationally equivalent data sets have been documented.

The universe of the data analyst is anchored in collections of raw observations. Data may be formatted to facilitate analysis and promote insight; however, from the analyst's perspective, the selection of a data format is a free and hence arbitrary choice, not itself a datum. The analyst's methods are thus constrained to treat identical observation sets equivalently, without regard to representational form. In this way, data analysis can be seen to embody an elementary but essential information concept – the signals are data representations, the world-states are data sets, and the information function maps each data representation to the data set or collection of data sets it accurately describes. This is the representation's "observational information content". Any two figures, tables, lists, matrices, or taxonomies are "observationally equivalent" if they accurately represent the same observation sets.

Methods of data analysis yield identical descriptions of observationally equivalent data representations, but human subjects do not. In *covariation assessment* tasks (reviewed in Allan, 1993; McKenzie, 1994), subjects are presented with an organized set of observations. Each observation is classified according to whether each of two categorical variables is present or absent, and the subject is asked to judge the strength of the relationship between the variables. A *joint-presence bias* in human covariation assessment has consistently been observed: judgments of relationship strength are more strongly impacted by joint-presence than joint-absence observations (e.g., Kao & Wasserman, 1993; Levin, Wasserman, & Kao, 1993; Lipe, 1990; Schustack & Sternberg, 1981; Wasserman, Dornier, & Kao, 1990). However, data can always be recoded so that the absence of an old variable is equivalent to the presence of a new variable. Therefore, any data representation with a high rate of joint presence is observationally equivalent to another data representation with a high rate of joint absence.

The observational equivalence of different presence/absence representations, and the consequences of a joint-presence bias, may be illustrated by a simple example. Suppose that every individual in a population is known to possess either Gene X or Gene Y, and also to possess either Trait 1 or Trait 2. Then the same sample of gene-trait observations from the population can accurately be represented by either of the matrix forms in Figure 2.1. Because of the well-documented empirical preference for joint-presence observations, human subjects are likely to infer a stronger gene-trait relationship when given the data summary in Form A rather than Form B. Observationally equivalent matrices thus lead to different judgments of relationship strength, in apparent violation of a Level 4 invariance requirement.

		Trait 1		Trait 2	
		Present	Absent	Present	Absent
Form A	Gene X				
	Absent	5	3	5	9
Form B	Gene Y				
	Absent	5	3	5	9

FIGURE 2.1 Observationally equivalent data summaries with different presence-absence coding.

However, the subject receiving a formatted matrix in a covariation assessment task is in a very different position from the analyst who elects to sculpt an initially formless mass of data into a convenient format. For the subject, but not for the analyst, the format is a datum – a signal received in the experimental environment. If it is a relevant datum that is, if the experimenter's selection of a data format leaks information relevant to statistical judgment – then matrices equivalent at Level 4 may not be equivalent at Level 1.

McKenzie and Mikkelsen (2007) argued that typical presence/absence coding leaks statistically relevant information about the relative rarity of variables. Human speakers, as a matter of cognitive and linguistic economy, tend to conceive and describe situations in terms of unusual properties, rather than the features they share with most events (see McKenzie, Ferreira, Mikkelsen, McDermott, & Strable, 2001, for experimental evidence bearing on this point). For the variables in terms of which human speakers typically code presence and absence, presence is thus likely to be ecologically rarer than absence. In this way, the experimenter's selection of Form A in Figure 2.1 provides evidence that Gene X and Trait 1 are relatively rare in the population at large.

Furthermore, observations of jointly rare variables have greater normative impact in covariation judgment, at least from a Bayesian perspective. To see why, consider a proposed statistical association between polydactyly (possessing more than 10 fingers) and dwarfism. Because both dwarfism and polydactyly are rare, it would not be unusual to observe a 10-fingered non-dwarf, regardless of whether dwarfism and polydactyly are correlated. However, although observing an 11-fingered dwarf would be unusual even if the variables were associated, it would be *extremely* unusual if they were not. If the presence of properties is rare, an observation of joint presence supplies stronger evidence for their covariation than an observation of joint absence. For formal details, see McKenzie (2004; McKenzie & Amin, 2002; McKenzie & Mikkelsen, 2000, 2007; see also Anderson, 1990; Oaksford & Chater, 1994).

Now consider the position of a hypothetical rational actor participating in a covariation assessment experiment. The actor receives limited information about a collection of observations involving unfamiliar variables with unknown natural frequencies. Among the limited information available to the actor is the presence

absence formatting of the observation summary. Because situations are normally described in terms of the presence or absence of rare variables, this suggests a lower prior probability, and correspondingly higher evidentiary impact, for joint-presence observations. Assuming the covariation assessment task is approached as a problem of statistical inference from the data set, rather than simple description of the set (McKenzie & Mikkelsen, 2007; Oaksford & Chater, 1994), we would expect the rational actor to exhibit a qualitative joint-presence bias. Notably, when subjects are explicitly informed that presence is common, the joint-presence bias weakens, and when subjects have prior knowledge that the presence of familiar variables is common, the joint-presence bias reverses (McKenzie & Mikkelsen, 2007), suggesting that the information analysis of presence-absence coding is relevant to the psychological analysis of its effects (cf. later "Two Views of Two Systems" section).

Level 5: Content Analysis

The levels of analysis reviewed above involve formally well-defined information concepts. Information is determined by Bayesian inference at Level 1, by semantic entailment at Level 2, by probability-outcome representation at Level 3, and by observational extension at Level 4. However, most human reasoning is informal, and separating informative wheat from rhetorical chaff is often more art than science. Individual domains of inquiry thus invoke rough norms of content analysis to distinguish relevant from irrelevant factors in the analysis of arguments. Because different rules of relevance apply in different domains, and because these rules are rarely explicitly formalized, the pragmatic norms of Level 5 comprise a family of information concepts that can only be loosely characterized.

For concreteness, we focus on a particular domain – political persuasion – in which framing phenomena have been widely studied (Chong & Druckman, 2007), which political life is awash with attempts at persuasion, populated by actors who devote elaborate attention not only to questions of "what to say" but equally to questions of "how to say it". In discussing studies of framing in political discourse, we adopt Druckman's (2001) useful distinction between "equivalency framing effects" and "emphasis framing effects". Equivalency framing refers essentially to the conventional effects of Levels 2 and 3 – that is, to experiments in which logically or formally equivalent frames happen to describe policy options (e.g., Quattrone & Tversky, 1988). Emphasis framing is more distinctive to the complex environment of political choice, and presupposes a Level 5 information concept, outlined below.

Emphasis frames typically are complex communications that defy the neat representational forms of Levels 2–4. For example, the Rose Garden setting of a President's campaign appearance may be viewed as setting an "incumbent frame" (Popkin, 1994). In the laboratory, experimental frames are often extended descriptions of issues or events that differentially emphasize competing values. In some cases, emphasis frames simulate media presentations – for example, the frames may be fictional news stories with headlines and photographs, or actual news reports from local television stations, reporting the same political event in different

ways (Nelson, Clawson, & Oxley, 1997a). Emphasis frames are considered equivalent if they influence preferences by differentially highlighting information of which the subject is already aware, rather than by supplying new information. Conceptual discussions of emphasis framing (Chong & Druckman, 2007; Nelson, Oxley, & Clawson, 1997b) sometimes assume that the frame recipient's attitudes can be captured, in some way, by an expectancy-value representation (Ajzen & Fishbein, 1980), in which attitudes are modeled as weighted sums of pre-evaluated attribute dimensions. For example, attitudes toward a government surveillance policy may be conceived as composite evaluations over such separable "dimensions" as cost, security, and liberty. Assuming such a representation, a communication is an emphasis frame if it changes the weighting of dimensions without changing the valuation of the policy on any particular dimension. In the context of this idealized model of beliefs and attitudes, the Level 5 information concept relevant to the analysis of emphasis frames can be approximated as described below.

The signals are communications delivered to a particular audience from a specific source. The world-states are possible policy attitudes of the audience, where these attitudes are assumed to be captured in a form akin to an expectancy-value representation. The information function maps a persuasive communication to those expectancy-value attitudes consistent with the truth of its overt factual claims and with the audience's prior beliefs. Two communications are "substantively equivalent" if they leave the audience with the same store of factual beliefs and attribute evaluations in memory, whatever the weights attached to them. To be sure, this characterization is little more than a hopeful sketch: expectancy-value representations of attitudes are operationally unavailable in practice, and may be psychologically inadequate in principle. In practice, then, judgments of substantive equivalence ultimately fall to researchers' intuitions, occasionally with limited corroboration from questionnaires probing subjects' interpretations of the frames – substantive equivalence generally cannot be decided by any simple algorithm (Sniderman & Therault, 2004). Accordingly, some framing researchers will reasonably question the value of this and similar Level 5 information concepts – that is, of attempts to informally equate and compare complex "frames" in special domains of discourse.

The often large effects of emphasis frames have sometimes evoked strong normative reactions. For example, Entman (1993, p. 57) suggested that political framing effects may "raise radical doubts about democracy itself. If, by shaping frames, elites can determine the major manifestations of 'true' public opinion that are available to government via polls or voting, what can true public opinion be? How can even sincere democratic representatives respond correctly to public opinion when empirical evidence of it appears to be so malleable, so vulnerable to framing effects?" However, the murky information analysis of emphasis frames complicates the normative assessment of their effects. Emphasis frames typically convey different information, explicitly (i.e., at Level 2) and/or implicitly (i.e., at Level 1). Researchers rarely attempt to explicitly characterize information content at these levels, assuming instead that bits of information distinctive to each frame are already subsumed in subjects' store of prior knowledge. A long tradition of research documenting citizen ignorance about policy and politics (Converse, 1964;

Delli Carpini & Keeter, 1996) suggests that such prior knowledge assumptions may be problematic. In some studies (e.g., Nelson et al., 1997b), researchers seek to rule out effects of new information by assessing correlations between respondent knowledge and frame susceptibility; however, these analyses often assume an oversimplified monotonic relationship between prior knowledge and posterior information, and their outcomes have been inconsistent across studies (Chong & Druckman, 2007).

While the information analysis of emphasis frames is necessarily inexact, it can be approximated in three stages: (1) a rough accounting of information, at Levels 1 and 2, that the frames differentially convey; (2) an assessment of whether this information is already known by all, or virtually all, subjects; and (3) when there may be an excess of (1) information over (2) prior knowledge for some subjects, a judgment is made as to whether this excess information could be relevant to their evaluations.

Consider, for example, an influential study by Nelson et al. (1997a) in which subjects were more tolerant of a specific Ku Klux Klan rally when the media report that described it (the "frame") consistently highlighted first amendment issues than when it consistently highlighted safety concerns. The safety frame may leak Level 1 information: (1) if media reports are more likely to highlight security concerns when violence is expected by experts, then recipients of the security frame may reasonably draw implicit inferences about the likely volatility of the event. Because (2) subjects would be unlikely to have detailed prior knowledge about the rally, (3) the selection of a media frame potentially contributes relevant information to its evaluation.

In framing effects involving evaluations of general policies, rather than specific events, the information analysis will depend on subjects' prior political knowledge, which will vary across populations. For example, in Nelson and Kinder's (1996) study of attitudes toward AIDS funding, a background blurb either characterized funding opponents as believing "that most people who get AIDS—primarily homosexual men and intravenous drug users—should have been more careful in the first place", or instead as believing "that the government has more important things to spend money on, like cancer research" (p. 1065). When respondents were subsequently asked for their personal views about AIDS funding, their opinions more strongly correlated with their attitudes towards homosexuals when they had received the "homosexual men and intravenous drug users" frame. These frames (1) plainly differ in their Level 2 information content. However, it is impossible to determine (3) the relevance of this information without knowing more about (2) the depth of the respondents' prior knowledge. For example, if respondents receiving the "cancer research" frame substantially updated their beliefs about the large-scale budgetary implications of increased AIDS funding, attitudes may legitimately have come to depend more directly on independent assessments of a range of alternative budgetary priorities, and hence less directly on views of homosexuals in particular. When different emphasis frames convey relevantly different information at Level 1 and/or 2, a psychological analysis may consider the possibility that their effects stem, at least in part, from the inferences they permit. The plausibility of such an analysis will commonly hinge on a careful

study of the extent of knowledge and the depth of ignorance in the subject population.

The Complexity of Information

The foregoing discussion of the levels of information illuminates the unexpected complexity of information analysis in typical framing tasks. Frames submitted to the equivalence method are designed to convey "the same information" in different ways, where information content is defined at a traditional level of normative analysis: The frames may be logically equivalent descriptions (Level 2), formally equivalent gambles (Level 3), observationally equivalent data digests (Level 4), or substantively equivalent attempts at persuasion (Level 5). But frames equivalent at Levels 2-5 are sometimes information non-equivalent at Level 1. As a consequence, the normative analysis of a standard framing effect is logically linked to the empirical analysis of natural regularities in frame selection. The next section considers experimental design strategies that would seek to exclude this empirical complexity from the information analysis of frames.

STRATEGIES OF INFORMATION ANALYSIS

Experimental frames are sometimes viewed as simple unmediated manipulations of the subject's point of view (a "frame in thought"), rather than as ambiguous utterances to which complex and largely non-conscious processes of interpretation will be applied, a "frame in communication"; Druckman, 2001. If it were possible to surgically implant frames into brains, so that particular descriptions of remembered choice scenarios would magically materialize in subjects' streams of thought, the regularities of conversational pragmatics could be excluded from the information analysis of frames. However, as surgical frame implantation is not on the horizon, are there more feasible methods for the systematic control of information leakage in experimental design? Two simple strategies are worth considering. Rather than eliminating the subject's representation of the external frame-selection process, an "information randomization" design would seek to control this representation and an "information elicitation" design would seek to measure it.

Unknown to subjects, the frame is a randomized experimental variable in standard designs. An information randomization design would overtly inform subjects, in some way, that frames are randomly generated (cf. Schwarz, Strack, Hillon, & Nadlerer, 1991). This step would seem to greatly simplify the normative analysis of effects: If the random nature of the signal-selection process can be firmly established, then the unrestricted information content of a frame no longer outstrips its logical, formal, observational, or substantive information content. If a well-established framing effect is lost with open randomization, this would seem to count as evidence that the original effect depended on implicit assumptions about frame selection. If the effect persists with open randomization, this would seem to count as evidence that the original effect did not.

By contrast, an information elicitation design attempts to rule out information leakage in the analysis of responses. After receiving a frame and indicating a choice, subjects are shown the alternative frame and asked whether, "on reflection", they did or would infer anything on receiving the first frame that they didn't or wouldn't infer had they instead received the second frame. Only those subjects who confidently affirm the equivalence of the frames are then included in the sample, to ensure that different responses to different frames are not mediated by different inferences. Indeed, Kahneman (2000b, p. xv) has argued that it is ultimately preferable to decide the equivalence of frames on the basis of subjects' reflective judgments about particular problems, rather than by appeal to a general formal theory that purports to be normatively binding: "It is the decision maker who should determine, after due consideration of both problems, whether the differences between them are sufficiently consequential to justify different choices. Violations of this lenient form of invariance demonstrate incoherence without a need for any judgment from on high about what is truly equivalent." A simple elicitation standard of equivalence has the advantage of circumventing questionable normative assumptions that general invariance principles may require, particularly at Level 3, cf. earlier "Levels of information" section.

However, empirical complexity is not easily cleared from the normative analysis of framing. Both information randomization and information elicitation have major shortcomings that narrowly limit their likely contribution to normative analysis. Both designs implicitly depend on simple – and, as we now show, contentious – psychological assumptions about the mechanisms of framing.

TWO VIEWS OF TWO SYSTEMS

As noted in the Introduction, normative evaluations of framing effects may take two forms. An information analysis asks whether the frames carry identical information. A psychological analysis asks whether the cognitive processes causing the effect are appropriately sensitive to the information carried by the frames. The present section sketches two psychological analyses of framing, formulated in a simplified two-system language. These contrasting psychological perspectives lend themselves to contrasting normative assessments of the cognitive processes responsible for the effects of information non-equivalent frames. They also clearly expose the shortcomings of the experimental strategies described in the previous section.

A currently popular dichotomy divides human cognition into two processing systems – a rapid, parallel, intuitive, effortless "System 1", and a slow, serial, deliberate, effortful "System 2" (e.g., Evans, 2003; Kahneman & Frederick, 2002; Stanovich & West, 2000). This division is undoubtedly oversimplified. It has been noted that human cognition is fractionated into many semi-independent systems, not just two (Evans, 2008), and that, in discussions of human rationality, the concept of a cognitive system is often applied incoherently (Keren & Schul, 2009). Nonetheless, the two-system picture can sometimes serve as a useful thumbnail sketch of cognitive structure, and it will simplify the (admittedly speculative)

discussion that follows. We emphasize, however, that the viability of the two perspectives considered here, and of the analysis to which they give rise, ultimately depends on the possibility of their reformulation in a richer language that distinguishes multiple systems, and distinguishes each system from the various processes of which it is composed.

We begin with what may be called "the standard view" of the two systems. In this view, the systems embody a simple trade-off between ease and sophistication of processing. System 1, with its formidable advantages in efficiency and horsepower, rapidly implements rough, ready, and relatively insensitive heuristics that streamline but also bias human judgment and choice. These biases can be submitted to, and ultimately vetoed by, the more nuanced and reliable assessments of System 2, but only at a steep cost in time and effort. This view of the two systems suggests a simple two-pronged strategy for the analysis of errors, summarized by Kahneman and Frederick (2002, p. 52): "In the context of a dual-system view, errors of intuitive judgment raise two questions: 'What features of System 1 created the error?' and 'Why was the error not detected and corrected by System 2?'"

In an alternative view of the two systems, systematic errors may arise from misplaced reliance on either. A key contention of this alternative view is that information barriers between the two systems are selectively impermeable in both directions – some information available to System 1 may be normally unavailable to System 2, and vice versa. The more generous capacity of System 1's parallel processing endows it with sensitivity to large sums of subtle information that System 2, with its limited-capacity serial computations, must largely ignore. This supposition is corroborated by a large literature documenting effects of stimuli that fail to reach explicit awareness, influencing behavior beyond the scope of serial conscious deliberation (e.g., Simons, Hamblin, Warren, & Day, 2007). On the other hand, there are likely to be systematic constraints on the transfer of novel information from System 2 to System 1: the process of "internalizing" into System 1 may take time, and may resist or dilute significant abstractions encoded by System 2. This more symmetrical view of their mutual information barriers suggests an alternative strategy for analyzing disagreements between the two systems. This strategy asks: Could relevant System 1 information have been crowded out of System 2's small-capacity representation? Alternatively, might relevant System 2 information have been incompletely internalized by System 1?

A second contention of this alternative view is that the two systems do not uniformly differ in sophistication of processing. The great long-term advantage of System 2 is its malleability over time. System 2 allows explicit principles, although often incorrect ones, to be formulated and submitted to argument and experiment (and thereby painstakingly and incrementally improved), and to be packaged for communication to other cognitive agents working on similar problems. These procedures open up tremendous vistas for System 2 in the long run. However, in facing a specific problem at a particular moment, System 2's principles, evolving slowly in the confines of a small roving spotlight of attention, may be less sophisticated than those implicit in the concurrent operations of System 1. System 2 is thus like the tortoise who is always behind but eventually wins. On this view, in the analysis of a specific disagreement between the two systems at an arbitrary stage in

System 2's history, no generalized assumption about relative sophistication of processing can be made – although over time the principles of System 2 may be expected to appreciate indefinitely in subtlety and power.

To be sure, these two views of the two systems are perspectives rather than testable hypotheses; any value they have lies in the economy with which they organize experimental observations, and the degree to which they ultimately lend themselves to theoretical refinements that generate specific predictions. However, these perspectives have traditionally drawn moral support from characteristically different kinds of experimental evidence. Arguments for the standard view typically point to studies involving deliberately idealized choice environments and employing relatively tight normative metrics – a research strategy typified by the equivalence method, as well as by apparently simple problems of abstract judgment (e.g., Wason, 1968). Arguments for the alternative view appeal instead to experiments placing subjects in richer choice environments regulated by less clear-cut norms. For example, recent experiments by Dijksterhuis and colleagues provide suggestive evidence for a putative “unconscious thought” process that seems better suited to computing over the multiple attributes of multiple complex choice alternatives (Dijksterhuis, 2004; Dijksterhuis, Baas, Nordgren, & van Baaren, 2006; Dijksterhuis & Nordgren, 2006) – though the interpretation and reliability of these experiments have been called into question (see Payne, Samper, Bettman, & Luce, 2008; and the meta-analysis in Acker, 2008). Perhaps relatedly, providing explicit reasons for choice sometimes leads to inferior hedonic outcomes in real situations (Wilson, Lisle, Schooler, Hodgger, Klaaren, & La Fleur, 1993). To be sure, each view affords alternative perspectives on the evidence usually advanced to bolster the other. Advocates of the standard view can reasonably criticize the loose normative metrics that more naturalistic choice experiments often require: The apparent normative superiority of System 1 in some studies may simply be an artifact of inadequate norms (e.g., Rey, Goldstein, & Perriechet, 2009). On the other hand, advocates of the alternative view can point to the hidden complexity that lurks behind many ostensibly simple judgment and choice tasks (e.g., McKenzie, 2003, 2005; Oaksford & Chater, 1994), and argue that the streamlined normative metrics that researchers use to analyze these tasks are System 2 oversimplifications, ignoring subtle information in System 1's vast store.

The equivalence method is a case in point. The information concepts of Levels 2–5 are distinctive inventions of System 2; they reduce the information content of a communication to a compact digestible kernel captured, at Levels 2–4, by a convenient formal principle, and they perform critical functions in the conduct of deliberate inquiry. When frames in a framing study are potentially information non-equivalent at Level 1, the two views of the two systems naturally favor two different interpretations of observed effects.

In the standard view of the two systems, the Level 1 analysis is seen as an esoteric System 2 refinement of the more conventional information analysis, even more remote from System 1's operations than the simpler System 2 information concepts of Levels 2–5. This refined information analysis may provide a fortuitous post hoc System 2 rationalization for the effect, but have nothing to do with the System 1 processes that actually generate it.

The alternative view of the two systems allows for this possibility, but seriously considers another. In the alternative view, Level 1 can plausibly be seen not as an esoteric System 2 refinement, but as the natural information level for the operations of System 1, with its wide capacity for processing subtle cues and its adaptation to ecological contingencies. Subtle information at Level 1, routinely exploited by System 1, may eventually be brought into a System 2 analysis of frames, but doing so requires a slow, serial, deliberate, and painstaking process of experiment and argument. The process of “externalization” from System 1 to System 2 is as uncertain as the process of “internalization” from System 2 into System 1.

The relative viability of these psychological assessments of a framing effect depends on multiple factors. How well-established are the proposed regularities in frame selection? Just how relevant is any information leaked by the different frames likely to be? And importantly, outside of traditional framing experiments, do subjects behave more generally in ways that are appropriately sensitive to the Level 1 leakage identified in the information analysis? To the extent that this is the case, the information analysis of frames is likely to translate into a strong psychological explanation of their effects. For example, because reference points have been found to affect speakers' selection of attribute descriptions across multiple content domains, and because listeners appear to draw implicit inferences about reference points both in traditional framing tasks and in other communication settings (McKenzie & Nelson, 2003; Sher & McKenzie, 2006), implicit sensitivity to Level 1 information is likely to be an important aspect of the psychology of attribute framing.

Alternatively, are specific psychological models for the effect available that (1) assume non- or counter-normative processing operations, (2) do not assume sensitivity to subtleties in the linguistic environment, and (3) successfully explain other effects arising outside of the communication environment in question? To the extent that this is the case, the information analysis of frames is less likely to contribute substantially to the psychological analysis of effects. For example, as noted above, some suggestive evidence of information leakage in the Asian disease problem has been reported (Sher & McKenzie, 2008; van Buiten & Keren, 2009). However, because this effect is predicted by prospect theory (Kahneman & Tversky, 1979) – which does not refer to regularities in the communication environment, and which purports to explain other phenomena occurring outside of the frame-selection environment (and possibly in other species; Chen, Lakshminarayanan, & Santos, 2006) – sensitivity to leaked Level 1 information may play only a supporting role in the psychological analysis of this problem.

A different approach to adjudicating between competing psychological perspectives on framing would adopt a correspondence rather than a coherence criterion of rationality (cf. earlier “Equivalence method” section). This approach would ask: Does susceptibility to framing make people better or worse off, overall, in normal social environments? While the standard view of framing effects would suggest that listeners should systematically suffer in some way from being “framed”, an “information leakage” perspective suggests that sensitivity to frames should make listeners better off overall in standard communication environments. While largely neglected in framing research, serious attempts to measure

the practical consequences of frame susceptibility in natural environments may be an important future avenue for indirect but valuable insights into the mechanisms responsible for framing effects. Notice that this ecological, correspondence-based approach reverses the standard rationale for the equivalence method, which was devised to substitute a conceptually clean coherence-based criterion for the messy complexity of calculating choice optima in natural environments.

INFORMATION STRATEGIES REVISITED

Earlier we considered the prospect of special experimental strategies that might simplify the information analysis of frames. An information randomization design would publicize the randomness of the frame-selection process, while an information elicitation design would ask subjects for their own judgments of equivalence. The shortcomings of these design strategies can be readily appreciated in the context of the two views of the two systems outlined above.

Consider first the information randomization design. Suppose that, when frame selection is openly scrambled, a well-established framing effect disappears. Would this show that the original effect stemmed from implicit assumptions about non-random frame selection? Not necessarily, according to the traditional view of the two systems. The explicit instruction about frame randomization may simply call conscious attention to the arbitrariness of the frames. As a consequence, System 2 may be unleashed on surface features of the frames that otherwise (i.e., in the absence of special attention) would trigger the simplistic heuristics of System 1. The standard conception of well-established framing effects could account for their disappearance with open randomization.

On the other hand, suppose that, when frames are openly randomized, a well-established framing effect persists. Would this show that normal regularities in frame selection were irrelevant to the initial effect? Not necessarily, according to the alternative view of the two systems. This view is non-committal with respect to the exact conditions under which warnings about special exceptions, delivered in abstract language to System 2, will be effectively "internalized" into System 1. The original effect may have reflected a normative response to the frames, on the basis of a generally accurate implicit default assumption, with the new effect simply reflecting a counternormative failure to update the adaptive default. The alternative conception of well-established framing effects could account for their persistence with open randomization.

Now consider the information elicitation design. Suppose that, when information judgments "on reflector" are directly elicited, subjects widely endorse the equivalence of the frames. As Kahneman (2000b) noted, an elicitation standard of frame equivalence has the virtue of doing away with excess theoretical baggage from formal normative models. Furthermore, the elicitation standard is comfortably at home in the standard view of the two systems. If one assumes that System 2 has ready access to System 1's full store of information, then the elicitation of System 2's reflective assessment is the simplest method for establishing the effective information equivalence of the frames. However, the information elicitation

design is inadequate in the alternative view of the two systems. If one allows that subtle information processed by System 1 is sometimes excluded from System 2's sparse representation of the same situation, then there is no guarantee that System 2 has access to Level 1 information driving System 1's response. The elicitation design only establishes that the two systems disagree. It does not determine which system is right.

CONCLUSION

The analysis of framing developed in this chapter can be loosely summarized in a single sentence: Framing effects are rarely, if ever, the one-shot self-contained demonstrations of human irrationality they are sometimes believed to be. The normative evaluation of a framing effect always rests on an information analysis of frames, and often rests on a psychological analysis of the cognitive processes they trigger. This information analysis, we have argued, is not neatly isolatable from the empirical analysis of human communication. If speakers are implicitly sensitive to relevant background conditions in selecting among frames, decision makers may reasonably be sensitive to the speaker's choice of frame in selecting among alternatives. The equivalence method, designed to evade the intractable complexity of normative optima, is confronted instead with the full richness of human frame selection in conversational environments.

Corresponding to the multiple levels of analysis at which the information content of frames can be defined, the intuitive processes that mediate their effects can be theoretically situated at different levels of information. In one perspective, intuition is seen as an inveterate simplifier, trading nuance for speed in the delivery of rough and ready assessments, equally remote from the subtle cues of Level 1 and the coarser signals of Levels 2-5. In an alternative perspective, the wide capacity of intuitive processing places it comfortably amid the complex ecological contingencies of Level 1, while serial conscious thought, with its sharp capacity constraints, plods slowly down the narrow corridors of Levels 2-5. Framing effects in complex information environments may thus be viewed through the lens of widely contrasting views of everyday rationality — as exemplifying the primitive simplicity, or the subtle sophistication, of human intuition.

NOTES

1. The probability representation of subjective uncertainty is treated here as a formal simplifying assumption. We do not consider the existence of such a representation of states of uncertainty itself to be a normative requirement although this is often contended; e.g., Savage, 1954. If probability representations are ultimately inadequate to the task of capturing actual human belief states, the Level 1 analysis of information may need to be correspondingly generalized.
2. Interestingly, however, an internal signal-selection process would thereby replace the usual external process of verbal communication, from the subject's perspective. Because memory selects its signals systematically, the form that a memory assumes

may permit reasonable meta-cognitive inferences that go beyond the explicitly remembered information.

3. The widespread use of frames in advertising may seem, *prima facie*, to suggest that framing effects are maladaptive, in that they can leave consumers susceptible to strategic manipulation by interested parties. However, this simple argument is problematic, for two reasons. First, every important system of information transmission opens potential avenues for deception, and hence new possibilities for strategic manipulation. In general, the possibility of strategic manipulation in communication reduces, but does not eliminate, the useful information content of signals, whether at Level 1 or at Levels 2-5. For rational information processors, the cost of being misled by "bad" communications will simply be outweighed by the benefits of inference from "good" communications. Second, the degree to which listeners can adaptively discount the Level 1 information content of frames as a function of a speaker's known persuasive agenda is an empirical question that has not been definitively addressed (cf. McKenzie & Sher, 2009).

REFERENCES

- Acker, F. (2008). New findings on unconscious versus conscious thought in decision making: additional empirical data and meta-analysis. *Judgment and Decision Making*, 3, 292-303.
- Allan, L. G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin*, 114, 325-348.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Azjen, I., & Fishbein, M. (1980). *Understanding attitudes and predicting social behavior*. Englewood Cliffs, NJ: Prentice-Hall.
- Chang, R. (1997). *Incommensurability, incomparability, and practical reason*. Cambridge, MA: Harvard University Press.
- Chen, M. K., Lakshminarayanan, V., & Santos, L. R. (2006). How basic are behavioral biases? Evidence from capuchin monkey trading behavior. *Journal of Political Economy*, 114, 517-537.
- Chong, D., & Druckman, J. N. (2007). Framing theory. *Annual Review of Political Science*, 10, 103-126.
- Converse, P. E. (1964). The nature of belief systems in mass publics. In D. Apter (Ed.), *Ideology and discontent* (pp. 206-261). New York: Free Press.
- Delli Carpini, M. X., & Keeter, S. (1996). *What Americans know about politics and why it matters*. New Haven: Yale University Press.
- Dijksterhuis, A. (2004). Think different: The merits of unconscious thought in preference development and decision making. *Journal of Personality and Social Psychology*, 87, 586-598.
- Dijksterhuis, A., Bos, M. W., Nordgren, L. F., & van Baaren, R. B. (2006). On making the right choice: The deliberation-without-attention effect. *Science*, 311, 1005-1007.
- Dijksterhuis, A., & Nordgren, L. F. (2006). A theory of unconscious thought. *Perspectives on Psychological Science*, 1, 95-109.
- Druckman, J. N. (2001). The implications of framing effects for citizen competence. *Political Behavior*, 23, 225-256.
- Eatman, R. M. (1993). Framing: Toward clarification of a fractured paradigm. *Journal of Communication*, 43, 51-58.
- Evans, J. St. B. T. (2003). In two minds: Dual-process accounts of reasoning. *Trends in Cognitive Sciences*, 7, 454-459.
- Evans, J. St. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255-278.
- Frisch, D. (1993). Reasons for framing effects. *Organizational Behavior and Human Decision Processes*, 54, 399-429.
- Gilovich, T. (1991). *How we know what isn't so: The fallibility of human reason in everyday life*. New York: Free Press.
- Gilovich, T., Griffin, D., & Kahneman, D. (2002). *Heuristics and biases: The psychology of intuitive judgment*. Cambridge: Cambridge University Press.
- Johnson, E. J., & Goldstein, D. (2003). Do defaults save lives? *Science*, 302, 1338-1339.
- Johnson-Laird, P. N., & Shafir, E. (1993). The interaction between reasoning and decision making: An introduction. *Cognition*, 49, 1-9.
- Jou, J., Shanteau, J., & Harris, R. J. (1996). An information processing view of framing effects: The role of causal schemas in decision making. *Memory and Cognition*, 24, 1-15.
- Kahneman, D. (2000a). A psychological point of view: Violations of rational rules as a diagnostic of mental processes. *Behavioral and Brain Sciences*, 23, 681-683.
- Kahneman, D. (2000b). Preface. In D. Kahneman & A. Tversky (Eds.), *Choices, values, and frames* (pp. ix-xvii). Cambridge: Cambridge University Press.
- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In Gilovich, T., Griffin, D., & Kahneman, D. (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 49-81). Cambridge, UK: Cambridge University Press.
- Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263-291.
- Kahneman, D., & Tversky, A. (2000). *Choices, values, and frames*. Cambridge, UK: Cambridge University Press.
- Kao, S.-F., & Wasserman, E. A. (1993). Assessment of an information integration account of contingency judgment with examination of subjective cell importance and method of information presentation. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 19, 1363-1386.
- Keren, G., & Schul, Y. (2009). Two is not always better than one: A critical evaluation of two-systems theories. *Perspectives on Psychological Science*, 4, 533-550.
- Kulbergner, A. (1998). The influence of framing on risky decisions: A meta-analysis. *Organizational Behavior and Human Decision Processes*, 75, 23-55.
- Levin, I. P. (1987). Associative effects of information framing. *Bulletin of the Psychonomic Society*, 25, 85-86.
- Levin, I. P., & Gaeth, G. J. (1985). How consumers are affected by the framing of attribute information before and after consuming the product. *Journal of Consumer Research*, 15, 374-379.
- Levin, I. P., Schneider, S. L., & Gaeth, G. J. (1998). All frames are not created equal: A typology and critical analysis of framing effects. *Organizational Behavior and Human Decision Processes*, 76, 149-188.
- Levin, I. P., Schmitzler, S. K., & Thee, S. L. (1988). Information framing effects in social and personal decisions. *Journal of Experimental Social Psychology*, 24, 520-529.
- Levin, I. P., Wasserman, E. A., & Kao, S.-F. (1993). Multiple methods for examining biased information use in contingency judgments. *Organizational Behavior and Human Decision Processes*, 55, 225-250.

- Lipe, M. G. (1990). A lens-model analysis of covariation research. *Journal of Behavioral Decision Making*, 3, 47-59.
- Mareau, T. M. (1989). Framing of information: Its influence upon decisions of doctors and patients. *British Journal of Social Psychology*, 28, 89-94.
- McKenzie, C. R. M. (1994). The accuracy of intuitive judgment strategies: Covariation assessment and Bayesian inference. *Cognitive Psychology*, 26, 209-239.
- McKenzie, C. R. M. (2003). Rational models as theories - not standards - of behavior. *Trends in Cognitive Sciences*, 7, 403-406.
- McKenzie, C. R. M. (2004). Framing effects in inference tasks - and why they are normatively defensible. *Memory and Cognition*, 32, 874-885.
- McKenzie, C. R. M. (2005). Judgment and decision making. In K. Lamberts & R. L. Goldstone (Eds.), *Handbook of cognition* (pp. 321-338). London: Sage Publications.
- McKenzie, C. R. M., & Amin, M. B. (2002). When wrong predictions provide more support than right ones. *Psychonomic Bulletin and Review*, 9, 821-828.
- McKenzie, C. R. M., Ferreira, V. S., Mikkelsen, L. A., McDermott, K. J., & Skrabble, R. P. (2001). Do conditional hypotheses target rare events? *Organizational Behavior and Human Decision Processes*, 85, 291-309.
- McKenzie, C. R. M., Liersch, M. J., & Finkelstein, S. R. (2006). Recommendations implicit in policy defaults. *Psychological Science*, 17, 414-420.
- McKenzie, C. R. M., & Mikkelsen, L. A. (2000). The psychological side of Hough's paradox of confirmation. *Psychonomic Bulletin and Review*, 7, 360-366.
- McKenzie, C. R. M., & Mikkelsen, L. A. (2007). A Bayesian view of covariation assessment. *Cognitive Psychology*, 51, 33-61.
- McKenzie, C. R. M., & Nelson, J. (2003). What a speaker's choice of frame reveals: Reference points, frame selection, and framing effects. *Psychonomic Bulletin and Review*, 10, 596-602.
- McKenzie, C. R. M., & Sher, S. (2009). *Product attribute framing and information leakage*. Unpublished manuscript.
- Nelson, T. E., Clawson, R. A., & Oxley, Z. M. (1997a). Media framing of a civil liberties conflict and its effect on tolerance. *American Political Science Review*, 91, 567-573.
- Nelson, T. E., & Kinder, D. R. (1996). Issue frames and group-centrism in American public opinion. *Journal of Politics*, 58, 1055-1078.
- Nelson, T. E., Oxley, Z. M., & Clawson, R. A. (1997b). Toward a psychology of framing effects. *Political Behavior*, 19, 221-246.
- Oakford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101, 609-631.
- Payne, J. W., Samper, A., Bettman, J. R., & Luce, M. F. (2005). Boundary conditions on unconscious thought in complex decision making. *Psychological Science*, 19, 1115-1123.
- Popkin, S. L. (1994). *The reasoning voter: Communication and persuasion in presidential campaigns*. Chicago: University of Chicago Press.
- Quattrone, G. A., & Tversky, A. (1985). Contrasting rational and psychological analyses of political choice. *American Political Science Review*, 82, 719-736.
- Rey, A., Goldstein, R. M., & Perruchet, P. (2009). Does unconscious thought improve complex decision making? *Psychological Research*, 73, 372-379.
- Rubinstein, A. (1995). *Making bounded rationality*. Cambridge, MA: MIT Press.
- Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.
- Schustack, M. W., & Sternberg, R. J. (1981). Evaluation of evidence in causal inference. *Journal of Experimental Psychology: General*, 110, 101-120.
- Schwartz, N., Strack, F., Hilton, D., & Nadler, G. (1991). Base rates, representativeness, and the logic of conversation: The contextual relevance of "irrelevant" information. *Social Cognition*, 15, 67-84.
- Shaler, G. (1958). Savage revisited. In D. E. Bell, H. Raiffa, & A. Tversky (Eds.), *Decision making: Descriptive, normative, and prescriptive interactions* (pp. 193-236). Cambridge: Cambridge University Press.
- Shafir, E. (1993). Choosing versus rejecting: Why some options are both better and worse than others. *Memory and Cognition*, 21, 546-556.
- Shafir, E., & LeBoeuf, R. A. (2002). Rationality. *Annual Review of Psychology*, 53, 491-517.
- Slavell, S. (2004). *Foundations of economic analysis of law*. Cambridge, MA: Belknap Press.
- Sher, S., & McKenzie, C. R. M. (2006). Information leakage from logically equivalent frames. *Cognition*, 101, 467-494.
- Sher, S., & McKenzie, C. R. M. (2008). Framing effects and rationality. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 79-96). Oxford: Oxford University Press.
- Simons, D. J., Hannula, D. E., Warren, D. E., & Day, S. W. (2007). Behavioral, neuroimaging, and neuropsychological approaches to implicit perception. In P. Zelazo, M. Moscovitch, & E. Thompson (Eds.), *Cambridge handbook of consciousness* (pp. 207-250). New York: Cambridge University Press.
- Sinderman, P. M., & Thierault, S. M. (2004). The structure of political argument and the logic of issue framing. In W. F. Sarris & P. M. Sinderman (Eds.), *Studies in Public Opinion: Attitudes, nonattitudes, measurement error, and change* (pp. 133-165). Princeton, NJ: Princeton University Press.
- Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences*, 23, 645-665.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211, 453-458.
- Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, 59, S251-278.
- Van Buiten, M., & Kerrin, G. (2009). Speaker-listener incompatibility: Joint and separate processing in risky choice framing. *Organizational Behavior and Human Decision Processes*, 108, 106-115.
- Wason, P. C. (1966). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, 20, 273-281.
- Wasserman, E. A., Dornier, W. W., & Kao, S.-F. (1990). Contributions of specific cell information to judgments of interevent contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 509-521.
- Williams, B. (1973). A critique of utilitarianism. In J. J. C. Smart & B. Williams (Eds.), *Utilitarianism: For and against* (pp. 77-150). Cambridge: Cambridge University Press.
- Wilson, D. K., Kaplan, R. M., & Schneidermann, L. J. (1957). Framing of decisions and selection of alternatives in health care. *Social Behavior*, 2, 51-59.
- Wilson, T. D., Lisle, D., Schooler, J. W., Hodges, S. D., Klaaren, K. J., & LaFlour, S. J. (1993). Introspecting about reasons can reduce post-choice satisfaction. *Personality and Social Psychology Bulletin*, 19, 331-339.