Framing effects on adolescent task representations, analytic and heuristic processing, and decision making Implications for the normative/descriptive gap

Paul A. Klaczynski*

Department of Psychology, 0417 Moore Building, Pennsylvania State University, State College, PA 16802, USA

Abstract

Adolescent decision making derives from the functioning of two independent information-processing systems. Analytic processing — the focus of most previous decision making and cognitive developmental research — often produces responses consistent with those that have traditionally been deemed correct. Heuristic processing, by contrast, often produces judgments and decisions at odds with existing norms. In part because heuristic processing is the system's "default," a large normative/descriptive gap exists. In the present research, early adolescents, middle adolescents, and young adults solved probability, contrary-to-fact, and sunk cost decision-making problems. The normative/descriptive gap decreased with age, but fallacious responses, assumed to reflect heuristic processing, were predominant across ages. "Framing" instructions to view the problems logically increased analytic processing and decreased, but by no means eliminated, the normative/descriptive gap. It is argued that by early adolescence, various heuristic processes and biases are firmly entrenched response tendencies that may represent overgeneralizations of strategies that are often adaptive. Prior to implementing decision-making interventions, the challenge of delineating conditions under which heuristic processing is adaptive and maladaptive must be met. © 2001 Elsevier Science Inc. All rights reserved.

Keywords: Decontextualization; Normative/descriptive gap; Representations; Rationality

1. Introduction

Research over the past three decades has demonstrated that people make astonishingly few judgments and decisions that are consistent with traditional normative standards. Adults

often misinterpret statistical base rates, overemphasize the value of vivid, readily available memories, make judgments based on misconstruals of probability data, violate the axioms of inductive and deductive inference, view the future unrealistically, select decision options prior to analyzing the consequences of those options, and change their performance standards on a moment-to-moment basis (Stanovich, 1999). Such large discrepancies between normative models and actual performance — the "normative/descriptive" gap (see Baron, 1985, 1988) — have led theorists down two closely associated paths, bound together by the mutual goal of explaining these puzzling findings.

The present paper is intended as an additional step down each of the trails currently being blazed. Specifically, a developmental approach was adopted to explore age-related differences in the normative/descriptive gap, relationships among normatively correct responses (NCRs) and biased responses on several judgment and decision making (JDM) tasks, and the effects of "frame" on judgments and decisions. Below, I outline each of the aforementioned theoretical paths and incorporate into this overview a description of two-process theories of decision making. Subsequent to this presentation, hypotheses derived from a two-process account of adolescent cognitive development are offered.

1.1. The normative status of JDM standards

The normative status of the guidelines (e.g., adherence to such statistical principles as Bayes' theorem) by which judgments and decisions are evaluated is the focus of rigorous debate (Stanovich & West, 2000; see Stanovich, 1999, for a comprehensive discussion of these debates). Common to these discussions are five types of argument. (1) It is the researchers themselves, rather than research participants, whose judgments of "correctness" are misguided (Cohen, 1981; Gigerenzer, 1996; Margolis, 1987); (2) the errors and biases observed in laboratory situations are random performance errors that reflect momentary lapses in the reasoning of otherwise rational beings (Stein, 1996); (3) participants' interpretations of task requirements deviate from those intended by researchers: People's judgments are usually rational, *given* the task construals on which they are based; (4) "Bounded rationality" theorists argue that poor judgments and decisions arise when information-processing capacities are overwhelmed. People are rational decision makers within biologically imposed limitations (Anderson, 1990; Oaksford & Chater, 1993; Simon, 1956); (5) Traditional norms are appropriate to many, but not all situations. The essence of this argument is that researchers who advocate traditional norms implicitly equate "optimality" with "precision" and "correct" processing with "exhaustive" processing (Reyna & Brainerd, 1995). Numerous goals other than precision, such as the desire to conserve cognitive energy, surround decision making; consequently, effective situational judgments are frequently misinterpreted as nonoptimal by existing normative guidelines. In short, general rules, principles, and axioms fail to capture the flexibility and complexity of everyday reasoning.

What can developmental research contribute to these discussions? To illustrate one possibility, consider argument (1): JDM errors are illusory and purportedly poor judgments are, in reality, adaptive (and thus should be considered "correct"). *If* this is the case, and *if* we assume (a) that cognitively advanced responding increases with age, then the

normative/descriptive gap *must* increase with age. Indeed, several investigators have documented precisely this phenomenon (Davidson, 1995; Jacobs & Potenza, 1991; Klaczynski & Narasimham, 1998; Reyna & Ellis, 1994). Either older children and adolescents are more irrational than younger children or the norms used to evaluate performance are incorrect. However, having accepted the proposition that rationality systematically covaries with age, it follows that the wrong standards (or standards that are inappropriate to some situations, the position (5) view) are being applied. Alternatively, consider the bounded rationality argument (4). If we assume (b) that information-processing efficiency increases through adolescence and culminates in early adulthood (e.g., Horn, 1970), then the normative/descriptive gap *must* grow smaller with increasing age (Klaczynski, 2000).

1.2. Two-process theories of judgments and decisions

Repeated observations of apparently irrational judgments and decisions have led to strenuous criticisms of traditional theoretical paradigms and have spawned numerous attempts to account for within-subject displays of normatively correct responding and of judgmental errors. Although questions regarding normative standards remain important, this path has been directed toward determining the cognitive and motivational processes under-lying decision making. The preponderance of these theories assert that decision making proceeds in a step-wise manner (e.g., encoding → retrieval of decision-making rules → cost/benefit analysis, etc.), hypothesize that reasoning frequently breaks down at one or more of these steps, and argue for a positive correlation between the extent of information processing and decision/judgment accuracy (i.e., "more processing is better processing"; see Reyna, Lloyd, & Brainerd, in press).[1]

Unlike the Piagetian, neo-Piagetian, information processing, and behavioral decision theoretical accounts, two-process theorists argue that no necessary relation exists between computational complexity and decision efficacy. Less extensive, less effortful nonanalytic proessing often produces optimal results. Also, unlike traditional explanations of cognitive development, two-process theorists do not subscribe to the notion that cognitive development progresses from predominantly intuitive reasoning during childhood to qualitatively superior, logico-mathematical reasoning in adolescence. Cognitive maturity does not entail the replacement of "intuitive" decision making by computational processing as the predominant means of interacting with the world (Reyna & Brainerd, 1995). Instead, two-process theories rely on the bifurcation of information processing into two independent, but interactive, systems. Although they function in qualitatively distinct ways, both a "heuristic" system and an "analytic" system are essential to cognitive adaptation. The co-

---

[1] For the sake of convenience only, "correct" responses are those judgments and decisions that correspond to those that have traditionally been prescribed as normative. Similarly, the term "error" does not necessarily connote maladaptive responding; rather, "error," "fallacies," and "biases" are simply deviations from widely accepted norms.

development tenet of two-process theories is, therefore, that the efficiency of both systems increases with age.[2]

This claim involves the tacit assumption that, because they are context-insensitive, traditional standards for assessing JDM quality are too inflexible to capture the fluid character of adaptive thought. To illustrate, consider the myriad situations that do not require precise formulations in order to make decisions. For example, in deciding whether to buy a half keg or a quarter keg of beer, an adolescent need not determine the number of people expected to attend a party, the number of drinks expected per attendee, or whether the extent to which the "Expected Attendance × Expected Average drinks/Attendee" product exceeds the number of drinks in a typical quarter keg. The judgment requires only a rough "more or less than a quarter keg" estimate: Calculating the precise number of drinks more (or less) than is provided in a quarter keg is cognitively wasteful, particularly given these uncertainties regarding multiplicands and the fact that kegs generally come in only two quantities, half and quarter. In such situations, decisions may reflect a trade-off between the motivation to conserve cognitive energy and the need for precision.[3]

Despite ongoing tension between these conflicting drives, heuristic processing is generally preferred to analytic processing. This preference arises, in part, because heuristic processing is cognitively economical, generally produces useful judgments, and is highly contextualized (Stanovich & West, 1998). These judgments are not consciously generated and thus seem to "pop" into consciousness (Sloman, 1996) and "feel" intuitively correct (Epstein, 1994).

Heuristic processing is contextually bound in the sense that unconsciously encoded situational factors (e.g., nonverbal cues), in combination with intrinsic factors (e.g., mood), often activate judgmental heuristics (e.g., "Go along with the crowd") that form the basis for consciously made decisions (e.g., "Buy half a keg of beer"; see Chen & Chaiken, 1999). A distinguishing aspect of heuristic processing, therefore, is that the judgment or decision that "comes to mind" (i.e., into working memory) is not the result of conscious efforts to reason through a situation or to retrieve a decision-making strategy. Consequently, the cognitive basis for such judgments is difficult to access and articulate. They may have the sense of insights, may arise from stereotype-based preconceptions and implicit memories, and may have a basis in various judgmental heuristics. Because heuristic processing requires little effort and occurs preconsciously, it frees cognitive resources for, and occurs in parallel with,

---

[2] Elsewhere, heuristic system processing has been referred to as associative, experiential, tacit, peripheral, intuitive, and implicit processing. Analytic processing has been referred to as rule-based, explicit, central, and rational processing (see Epstein, 1994; Evans & Over, 1996; Reber, Walkenfeld, & Hernstadt, 1991; Sloman, 1996; Stanovich, 1999).

[3] I use the term "motivation" loosely to include both conscious and unconscious drives toward goals that may also be either consciously or unconsciously determined. Also, note that the motivation to conserve cognitive energy does not imply that there exists a fixed reservoir of cognitive resources. In principle, cognitive resources may be infinite. The motivation to conserve effort may arise not because resources become increasingly sparse following cognitive expenditures (and are replenished only after some [unknown] amount of time), but instead because expending effort — particularly to support the goals of another person (e.g., the experimenter) — is inherently aversive (consider, for example, students' aversion to studying for examinations, particularly for courses that have no direct relevance to their goals).

computationally complex reasoning. Despite the possibility of parallel processing, however, there are cases in which the lure of ''quick and dirty'' heuristic processing almost entirely usurps computational reasoning (Evans, 1996).

In contrast, analytic processing is consciously controlled, effortful, and deliberate, and is dependent on the acquisition and utilization of abilities that are frequently prescribed as optimal for decision making and that have long been seen as the apex of cognitive development. Examples of analytic competencies include the higher-order abilities that enable the calculation of likelihood ratios, reasoning consistent with the dictates of propositional logic, judgments based on comparisons between a priori probabilities, and logically consistent reasoning across problems that differ in their superficial qualities. When analytic processing is predominant, individuals are metacognitively aware of engaging in decision making. Even so, heuristic processes may continue to affect decisions in ways that cannot be consciously controlled (Bargh & Chartland, 1999; Chen & Chaiken, 1999).

Heuristic and analytic processing may function independently, but nonetheless, may simultaneously contribute to the same task. During the initial stages of problem solving, parallel processing enables multiple task representations. Verbatim representations derive from detail-oriented processing and correspond closely to actual problem information. Such representations, which involve considerable effort to maintain in working memory, generally function to facilitate declarative memory. Gist representations, in contrast, are holistic abstractions that impute subjective meaning — which may or may not correspond to the meanings ascribed by others — into a task and provide frameworks around which solutions can be constructed (Reyna & Brainerd, 1995).

The distinction between verbatim and gist representations is critical because, in contrast to the assumptions of information-processing theories, verbatim representations are less central to optimal decision making than are gist representations (Reyna et al., in press). Whereas verbatim representations sometimes provide quantitative information for computational reasoning, gist representations are insights into the structural requirements of a task.

Logically isomorphic tasks may be represented differently because irrelevant aspects of one task (e.g., content elicits misleading beliefs and memories) interfere with gist abstraction, because accuracy motivation differs between tasks, and because the contexts that ''frame'' otherwise identical tasks differ (see Reyna et al., in press). Although interference arises from multiple sources, the thread tying these together is that they decrease the probability of extracting accurate, ''decontextualized'' task representations (decontextualization involves separating task content, and beliefs associated with that content, from underlying task structure; Stanovich & West, 1997). Attempts to abstract decontextualized representations do not, however, guarantee useful representations because, for example, variations in problem complexity, instructions, and demands for precision may mislead or overwhelm problem solvers and because heuristic interference is difficult entirely eliminate (Klaczynski & Gordon, 1996).

This analysis implies that the system predominant in a given decision-making situation is jointly determined by the type (i.e., contextualized, decontextualized) and quality (i.e., precise, imprecise; accurate, inaccurate) of gist representation and by motivation (Klaczynski, 2000).

Reasoners prefer to rely on the least precise gist representation that will yield a solution because such gists originate directly from heuristic processing and thus arise with little or no conscious effort. Such representations, in turn, "pull" reasoners toward "easy" (i.e., heuristic) cognitive operations. In contrast, more precise, decontextualized gists are avoided because these sometimes arise only after extensive analytical processing of verbatim information and because operating on these representations often requires additional analytic, cognitively expensive, resources.[4]

Consider the task of determining which of two lotteries is most likely to yield a "winner." In Lottery A, one winning ticket will be pulled from a total pool of 10 tickets. In Lottery B, there are 10 winning tickets in a pool of 100 tickets. "Denominator neglect" (DN) (exclusive focus on the size of numerators) is a common source of errors on such tasks because the reasoners rely on the low-level gist representation, "Compare numbers of winners," instead of the more precise, "Compare ratios" (see Reyna et al., in press). If the former representation is selected, no reasoning (in the usual sense of the term) is required for a judgment; processing is almost entirely heuristic and individuals simply choose Lottery B. If the latter representation is selected, then subsequent processing will involve effortful computations.

As this example suggests, problems, errors, and biases arise when the two systems "pull" for different solutions (Denes-Raj & Epstein, 1994; Stanovich, 1999) *and* (1) cues for heuristic processing overwhelm cues for analytic processing (e.g., highly contextualized tasks activate irrelevant beliefs and stereotypes), (2) problem content produces memory interference, thereby inhibiting analytic reasoning; (3) the analytic competence to solve a problem correctly is lacking, (4) the requisite competence has been acquired, but the wrong gist representation of a task is selected (e.g., an imprecise representation is used for a computational task, as in the example), (5) insufficient time is available for analytic processing, and (6) the relevant competence is activated, but attempts to implement that competence fail (see Reyna et al., in press, for discussion of these and other impediments to normative responding).

1.3. The present investigation

The purpose of the present study was, first, to explore the developmental trajectories of normatively correct reasoning and of judgmental fallacies across three disparate tasks. A second purpose was to investigate how instructions to "frame" the problems heuristically or analytically affect the normative/descriptive gap.

Because the quality of problem representations is influenced by the perceived need for precision, it may be possible to increase analytic processing, decrease heuristically generated fallacies and biases, and reduce the normative/descriptive gap by instituting minor changes in the phraseology of a task. That small differences in wording can produce large discrepancies

---

[4] Terms such as "preference" often connote conscious decisions; however, preferences may be unconsciously motivated. For example, when a teacher consistently writes on chalkboards in script rather than in print, he or she is indicating a preference behaviorally but may not have consciously decided from among available options.

in responding has been amply demonstrated in vast literature on "framing" effects (Frisch, 1993; Tversky & Kahneman, 1981; see, however, Fischhoff, 1988). Here, following Epstein, Lipson, Holstein, and Huh (1992), the initial question for each problem was intended to elicit the typical representation on which participants base their judgments and decisions. A second question was intended to elicit an analytic frame by merely requesting that participants view the problem from the perspective of a "perfectly logical person."[5]

The present study is an extension of prior research involving this framing manipulation, which has focused exclusively on college students, to early and middle adolescents. Findings that this subtle manipulation increase NCRs across diverse JDM tasks and across adolescent age groups would advance current theories of adolescent JDM by (1) indicating that the size of the normative/descriptive gap is more fragile than is often assumed and, therefore, is more open to emendation than bounded rationality theorists often presume (see Stanovich, 1999, pp. 234–236), (2) suggesting developmental invariance in the effects of the framing manipulation on task representations, and (3) supporting the hypothesis that representation quality mediates the relationship between processing system and performance. (4) The presence of systematic fallacies in the "typical" frame would indicate that JDM errors do not arise from "momentary lapses of reason," as some authors have argued (e.g., Stein, 1996). (5) If these fallacies and biases are pervasive and if they do not diminish with age, then arguments that the default processing system is heuristic (as opposed to analytic) and that the preeminence is this system does not decline during the adolescent years would be supported.

2. Method

2.1. Participants

Fifty-nine early ($n = 28$; mean age = 12.81, S.D. = 0.96 years) and middle ($n = 31$; mean age = 16.77, S.D. = 1.96 years) adolescent volunteers and 31 young adults (mean age = 21.74, S.D. = 3.28) participated. The adolescents were recruited from the 7th and 8th grades of two junior high schools and from the 10th through 12th grades of a local high school. All adolescents were drawn from advanced and/or college preparatory classes. Of the 59 adolescents, 49 indicated certainty that they would attend college; 10 (six junior high school; four high school) indicated uncertainty, but none indicated that they were definitely not going to college. The adolescents participated individually or in groups of two to three students. The

---

[5] The term "frame" is centered on the relationship between task content and task representations. Similar tasks are differently framed when semantic cues, embedded in either task content or question phrasing, lead to different representations. Its use here, however, is "loose" (see Frisch, 1993; Stanovich, 1999) relative to the stricter sense of the term implied by Kahneman and Tversky (1984). The framing manipulation used here is similar to, but not identical with, instructions for precise responding. Precision instructions, however, do not reduce or eliminate certain biases (e.g., overconfidence; Fischhoff, 1988). The discrepant findings of this research and of precision instruction research indicate some task specificity in the effectiveness of framing and precision instructions, a point to which I return in the Discussion.

young adults, recruited from introductory psychology classes, participated in groups of three to eight students.

2.2. Procedure

The goals of this research were met by using JDM tasks for which, among adults, there exist well-established response tendencies (see Baron, 1988; Denes-Raj & Epstein, 1994; Frisch, 1993). Specifically, although they present very different challenges, the tasks described subsequently were selected because each has been consistently associated with surprisingly low rates of NCRs and with specific biases (e.g., DN).

In a single 15–20-min session, participants were presented seven-page problem packets, which contained six JDM problems. Each packet contained a cover sheet, two probability problems, two counterfactual reasoning problems, and two sunk cost problems. The six problems were shuffled immediately prior to presentation to guard against order effects. Participants were instructed to read each problem carefully, to take their time, and to respond to each question.

After each of the six problems, two "framing" instructions were presented. The intent of one question (the "usual" question) was to elicit participants' default manner of processing. The purpose of the second question (the "logic" question) was to elicit a shift from participants' "usual" processing to analytic processing. For participants who responded analytically to the "usual" question, the "logic" question should not influence responses. Both questions were adapted from Epstein et al. (1992).

For each problem, the "usual" frame always began, "Think about this situation as you normally would." This instruction was followed by a question specific to the problem. The "logic" frame always began, "Think about this situation from the perspective of a *perfectly logical person*." This instruction was followed by the same problem-specific question that followed the "usual" question, although it also contained a reminder to respond logically.

Two scores were given for each problem; one point was given for each NCR and one point was given for each fallacious response. Note that these scores are not mutually exclusive because, although it was not possible to receive two scores of "1" on the same problem (i.e., a response could not be both normatively correct and fallacious), the typical fallacies for each task were not the only possible non-normative responses. Analyses of both scores were therefore conducted because they could yield different information regarding developmental trends and framing effects. NCR scores and fallacy scores were summed across the two problems of each JDM task and averaged (i.e., total NCR score/2; total fallacy score/2), yielding four scores for each task — proportions of NCRs and proportions of fallacious responses in the "usual" and "logic" questions.

2.3. Materials

2.3.1. Probability judgments

In each probability problem, two options were presented that had identical ratios of targets to totals [i.e., targets/(targets + nontargets)]. The only differences between options were the

numbers of targets (i.e., the numerator) and totals (i.e., the denominator). Participants decided which (if either) of the two options he or she would prefer. In one problem, the probability of enrolling in a course with an award-winning professor was 8% in both options (2/25 at one university; 4/50 at a different university). In the example below (based on Denes-Raj & Epstein, 1994), the probability of selecting a winning lottery ticket from either of two jars was 10%.

You are playing a lottery in which you can win $1000. There are two jars from which you can select a winning ticket. In the first jar, there are only 10 tickets, and 1 of these is the winning ticket. In the second jar, there are 100 tickets and 10 winning tickets.

Think about this situation as you normally would. Which jar, if either, would you select from to have a better chance of winning the lottery?

  a. The jar with 1 winning ticket
  b. The jar with 10 winning tickets
  c. It would not matter to me

The problem was then presented again, followed by instructions as mentioned below.

Think about this situation from the perspective of a *perfectly logical person*. Which jar, if either, would a perfectly logical person select from to have a better chance of winning the lottery?

(*Note*: option "c," the NCR, for the logic question was "It would not matter to him or her.")

Although the proportions of targets in options "a" and "b" are identical, prior research (see Denes-Raj & Epstein, 1994; Reyna et al., in press) has shown that both children and adults often fail to consider differences in denominators (hence, "DN") and favor the option ("b") with the greatest absolute number of targets.

2.3.2. Counterfactual "if-only" (IO) judgments

Generally in counterfactual thinking, individuals believe that a current state of affairs would nonetheless be the same even if preexisting conditions had been different from actual prior conditions (e.g., "I would still be a minister even if I had been raised by hedonistic parents."). A subtype of counterfactual thinking has been referred to as the IO fallacy (Epstein et al., 1992). The IO fallacy occurs when behaviors are judged more negatively when it *appears* that a negative consequence could have been easily anticipated, and therefore avoided, in one of two logically identical and equally unpredictable situations. Consider the example below (adapted from Epstein et al., 1992).

Tom parked his new car in a parking lot that was half empty. His wife asked him to park in a spot closer to where she wanted to shop, but he parked, instead, in a spot closer to where he wanted to shop. As luck would have it, when he backed out after shopping, the car behind him backed out at the same time, and both cars sustained about $1000 worth of damage.

Robert parked his car in the same parking lot when there was only one parking place, so he took it. As luck would have it, when he backed out after shopping, the car behind him backed out at the same time, and both cars sustained about $1000 worth of damage.

In both cases, the accidents were unpredictable and not under the control of the involved parties. Yet low-level gist representations (e.g., Tom had control, Robert had no control) activate heuristics that link control to fault (i.e., similar to the "fundamental attribution error," wherein observers overestimate the role of dispositional factors when assessing a person's actions). Tom, whose accident appeared as though it could have been avoided if he had heeded his wife, is believed by most young adults to have made a worse decision than Robert (Epstein et al., 1992).

For each IO problem, the scenario was presented twice, once followed by the "usual" question and once followed by the "logic" question. Participants indicated which, if either, of the two involved parties had acted "more foolishly" on five-point scales. In the example, the scale points were: "Robert acted much more foolishly than Tom," "Robert acted somewhat more foolishly than Tom," "Robert and Tom acted equally foolishly," "Tom acted somewhat more foolishly than Robert," and "Tom acted much more foolishly than Robert." From a normative perspective, the correct judgment is that the two parties were equally "foolish." The IO fallacy occurred only when Tom's decision was rated as worse than Robert's (rather than the reverse).

2.3.3. Sunk cost decisions

The sunk cost fallacy occurs whenever current judgments and decisions are influenced by inconsequential past decisions. Subsequent to investing in a goal and then discovering that the goal is no longer worthwhile, is no longer attainable, or is less important than competing goals, an individual "honors" sunk costs when he or she nonetheless continues investing in and pursuing the goal. Generally, the decision to honor sunk costs is justified by claims that resources in a goal are "wasted" if the goal is not pursued, as in the decision of a student who remains in a course because "I've already put three quarters of a semester into the class" despite having no possibility of passing (the remaining time could be better invested in other courses).

Two sunk cost problems (adapted from Frisch, 1993) were presented. Each problem described two situations that differed only in terms of whether or not costs had been sunk into a decision. See example given below.

Consider the following two situations and respond to the questions that follow each.

A. You are staying in a hotel room on vacation. You paid $10.95 to see a movie on pay TV. After 5 minutes, you are bored and the movie seems pretty bad. How much longer would you continue to watch the movie?

B. You are staying in a hotel room on vacation. You turn on the TV and there is a movie on. After 5 minutes, you are bored and the movie seems pretty bad. How much longer would you continue to watch the movie?

Because sunk costs (i.e., money paid for the movie) are irrelevant (i.e., they are irretrievable and should be ignored), decisions in the two situations should be the same. The task was to select, from among five options, a course of action in each situation. In the example, options for both Scenario A and Scenario B were: stop watching entirely, watch for

10 more minutes, watch for 20 more minutes, watch for 30 more minutes, and watch until the end. After making a decision for the "usual" question, each problem was presented again and participants responded to the "logic" question.

Responses were considered normatively correct when decisions in the two situations were identical. The decision to spend more time/effort in the "sunk cost" situation than in the "no sunk cost" situation indicated the sunk cost fallacy (the opposite response is also non-normative, but does not indicate the typical sunk cost fallacy).

3. Results

Mean proportions of NCRs and fallacies on each of the three JDM tasks (i.e., probability, IO, and sunk cost) for each age group and in each frame are presented in Table 1. Analyses of each task involved 3 (Age: early adolescent, middle adolescent, young adult) $\times$ 2 (Frame: usual, logic) analyses of variance.

Table 1
Mean proportions of normatively correct, fallacious, and biased judgments and decisions for each problem type and age group in the usual and logic frames

Problem type	Early adolescent	Middle adolescent	Young adult
Probability NCR			
Usual	0.07 (0.18)	0.19 (0.31)	0.21 (0.34)
Logic	0.23 (0.25)	0.39 (0.38)	0.42 (0.32)
Denominator neglect			
Usual	0.57 (0.33)	0.53 (0.34)	0.58 (0.37)
Logic	0.32 (0.37)	0.27 (0.28)	0.39 (0.31)
If-only NCR			
Usual	0.14 (0.23)	0.19 (0.28)	0.27 (0.39)
Logic	0.32 (0.37)	0.53 (0.39)	0.66 (0.42)
If-only fallacy			
Usual	0.54 (0.30)	0.63 (0.39)	0.63 (0.34)
Logic	0.32 (0.34)	0.39 (0.33)	0.27 (0.36)
Sunk cost NCR			
Usual	0.16 (0.31)	0.27 (0.36)	0.37 (0.36)
Logic	0.21 (0.25)	0.31 (0.38)	0.44 (0.36)
Sunk cost fallacy			
Usual	0.55 (0.39)	0.55 (0.40)	0.50 (0.29)
Logic	0.43 (0.30)	0.26 (0.31)	0.31 (0.31)

Standard deviations are in parentheses.

3.1. Probability judgments

3.1.1. Normatively correct responses

Both age and frame were significantly related to NCRs on the probability problems, $F(2,87) = 3.08$, $P = .051$, and $F(1,87) = 38.06$, $P < .001$, respectively. Across frames, young adults and middle adolescents were more likely than early adolescents to indicate that the absolute difference between numerators was irrelevant to their choices. Although the frequency of NCRs more than doubled from the usual frame to the logic frame, only 41% of the young adults, and 35% of the entire sample, made normative judgments in the logic frame. In the usual frame, NCRs were observed on only 16% of the problems. The low numbers of NCRs is particularly disconcerting given the apparent simplicity of the problems and that responding was well below chance (i.e., 33%). Participants were neither calculating ratios nor guessing: The significant deviation from chance responding indicates a systematic biasing influence.

3.1.2. Denominator neglect

As noted earlier, the "DN" bias arises when larger absolute numbers of successful acts (e.g., 10 winning tickets) are perceived as more probable than smaller absolute numbers of successful acts (one winning ticket), despite identical probabilities of success.

In the usual frame, examination of DN biases indicated that slightly better than 53% of responses in each age group chose the alternative with largest absolute number. A significant effect of frame, $F(1,87) = 53.25$, $P < .001$, was indicated by a drop in DN percentages to 33% in the logic frame. Age was not related to DN errors, $F < 1$, indicating that the previously reported developmental increase in NCRs was not paralleled by a decrease in DN bias. Instead, this null effect indicates that the age-related increase in NCRs was accompanied by fewer judgments that the smaller numerator option (e.g., one winning ticket,) was more probable than the larger numerator option (e.g., 10 winning tickets).

3.2. IO judgments

3.2.1. Normatively correct responses

Analyses of IO NCRs revealed significant effects of age group and frame, $F(2,87) = 5.22$, $P = .007$, and $F(1,87) = 51.99$, $P < .001$, respectively. Across frames, young adults made the NCR (neither party was more or less "foolish" than the other) more often than middle adolescents who, in turn, gave more NCRs than early adolescents. As with the probability problems, NCRs were higher for all age groups in the logic frame than in the usual frame.

Despite age and frame effects, normative responding was infrequent: In the usual frame, NCRs were given on less than 21% of the problems — the level expected by chance. Only middle adolescents and young adults, and only in the logic frame, gave NCRs on better than 50% of the problems.

3.2.2. The IO fallacy

The next analyses were conducted to examine age and frame effects on the IO fallacy, indicated by judgments that decisions associated with seemingly foreseeable and seemingly

controllable outcomes were worse than decisions whose outcomes were more evidently the results of bad luck. As with probability problems, fewer fallacies were committed in the logic frame than in the usual frame, $F(1,87) = 44.80$, $P < .001$. In neither frame were age differences in the IO fallacy significant, $F = 1.14$.

Similar to the probability judgment findings, developmental increases in NCRs on the IO problems were not accompanied by age-related declines in the IO fallacy. The NCR increase thus reflects a decline in "counter-IO" judgments (i.e., "bad luck" decisions related to apparently *un*predictable outcomes were judged as worse than decisions for which negative outcomes were apparently foreseeable).

3.3. Sunk costs

3.3.1. Normatively correct responses

Responses to the sunk cost problems were considered correct if the same decisions were made in the scenarios in which investments had and had not been made. Analyses of these responses revealed that sunk costs were ignored more often by adults and middle adolescents than by early adolescents, $F(2,87) = 4.00$, $P = .022$, and in the logic frame more often than in the usual frame, $F(1,87) = 6.82$, $P = .011$. Even in the logic frame, however, across age groups only 35% of the responses were normatively correct.

3.3.2. Honoring sunk costs

Fallacious responses, those which honored sunk costs, were indicated by decisions to continue poorly chosen actions because of prior investments in those actions. Across framing questions, the percentage of these responses was not age-related, $F < 1$. In the usual frame, at least 50% of the responses in each age group indicated sunk cost reasoning. Thus, instead of being paralleled by declines in sunk cost reasoning, age-related gains in NCRs were a function of age-related decreases in the tendency to pursue ill-fated plans in which *no* investments had been made. Finally, although the tendency to honor sunk costs was by no means eliminated, the fallacy was committed much less frequently in the logic frame than in the usual frame, $F(1,87) = 29.15$, $P < .001$.

4. Summary

Different analytic competencies were required to generate NCRs on the three JDM tasks studied here. Probability judgments require the calculation and comparison of ratios. IO judgments require understanding that the appearance of predictability may be illusory and that post hoc, "Monday-morning quarterback," explanations are fragile; the ease with which decision → subsequent event causal explanations can be generated after the fact is irrelevant to decision quality. Sunk cost decisions require dissociating current decisions from investments in a prior decision or course of action.

Despite these substantially different requirements, findings were similar across tasks. First, normatively correct responding increased with age. However, in the usual frame, less than

one third of the responses of the two older age groups were normative; surprisingly, *no* participant responded normatively to all six problems. Second, normative answers improved and fallacies decreased in the logic frame. It is unlikely that these effects occurred because the logic frame always followed the usual frame. In Epstein et al.'s (1992, Exp. 2) research and in a pilot study conducted by the author, order of question presentation did not impact responding. Nonetheless, individual differences in interpretations of the framing questions — some individuals, for example, may not have responded to the instructions because they assumed that their "usual" perspective *was* the "logical" perspective — may, to some extent, explain why NCRs remained infrequent in the logic frame. For the present analysis, however, it is important to note that variations in instruction interpretations were not related to age: The change in perspective from "usual" to "logical" was as effective for the early adolescents as it was for the older participants.

Third, in neither the usual nor the logic frame did the fallacies typically made on these tasks decline with age. This finding suggests that the heuristic "pull" of these problems does not diminish with age. In other words, as anticipated in the Introduction, the pervasiveness of systematic fallacies across tasks indicates that the problems elicited heuristic processing regardless of age.

Fourth, within-subject variability in responding — indicated by shifts from the typical fallacy on one problem to the NCR on the second problem of the same task (and vice versa) *and* by shifts from the typical fallacy on one problem to an atypical error on the second, logically isomorphic problem — may provide insight into the question, What develops? When not drawn to the typical fallacy, early adolescents frequently committed atypical errors. With increasing age, when the strong heuristic attraction of the fallacious response was avoided, responding was more analytical.

In Fig. 1, proportions of NCRs, "typical" fallacies, and "atypical" errors (i.e., errors prior research indicates are uncommon among adults), collapsed across framing conditions, are presented. At the outset of adolescence, variability seems to occur primarily within the heuristic processing system, at least on the tasks used here. In the figure, this is reflected by high proportions of typical fallacies and of atypical errors. Later in development, variability is more likely to arise from shifts between heuristic processing and analytic processing, reflected by the high rates of fallacies and NCRs and by comparatively few atypical errors.

Several explanations can be offered for the predominantly heuristic variability in the responding of the early adolescents: (1) The tasks were adequately represented, but the relevant normative principles had not yet been acquired. For reasons elaborated in the Discussion, this possibility is unlikely. Other possibilities are: (2) the early adolescents were aware of the relevant principles but did not accept them as superior to heuristic responses, and (3) although representations differed in the usual and logic frame, in both frames, these were inaccurate reflections of task demands and thus cued different heuristic responses. These speculations suggest that, with increasing experience, socially shared responses (i.e., "typical" fallacies) may become more deeply entrenched in memory and may be applied automatically to decisions for which they are sometimes ineffective (see Arkes & Ayton, 1999, for a similar argument concerning sunk cost judgments). Relatively idiosyncratic

Fig. 1. Proportions of NCRs, "typical" fallacies, and "atypical" errors, collapsed across framing conditions and across the probability, sunk cost, and IO tasks. Approximately 81%, 68%, and 61% of all responses were either typical fallacies or atypical errors among the early adolescents, middle adolescents, and young adults, respectively.

response tendencies may eventually drop out of adolescents' cognitive repertoires as a function of their inefficiency (see Brynes, 1998; Siegler, 1996).

5. Discussion

In this final section, I address the implications of the findings for "bounded rationality" views of JDM and for the misguided norms perspective, for discussions of the normative/descriptive gap, for two-process approaches to adolescent decision making, and for decision-making interventions.

5.1. Bounded rationality, computational limitations, and framing effects

Bounded rationality theorists (e.g., Evans & Over, 1996) have argued that biological constraints on information-processing capacity prevent most adults from reasoning according

to traditional normative principles. However, within the boundaries of these constraints, humans are rational decision makers. An important implication of this view is that the production of normative solutions is beyond the processing capacity of most individuals. It thus follows that these norms should not be the standards by which decisions are judged (Baron, 1985; Stanovich, 1999). This argument further implies that if the biological foundations of cognition are firmly established by early adolescence and if most adolescents lack the requisite capacity, programs intended to foster understanding, acceptance, and application of many normative principles are doomed to failure.

In this section, the questions, "Can the low rate of normatively correct responding across age groups be attributed to information-processing limitations?" and, "Can developmental increases in NCRs be attributed to increases in cognitive capacity?" are addressed. I do not argue that the findings are definitive on either score; yet, I would suggest that the probable answer to each question is "no."

Consider first that NCRs were least common on the probability problems, even though probabilistic understanding — specifically, the abilities to compute ratios and to make predictions derived from ratios — develops before early adolescence (Kreitler & Kreitler, 1986). It is similarly difficult to argue that the low rates of NCRs on the sunk cost problems are attributable to competence/capacity limitations. Specifically, existing research indicates that preadolescents avoid the sunk cost fallacy, but that adults do not (Arkes & Ayton, 1999). The notion of content interference — that logically irrelevant memories, strategies, and heuristics are elicited by task content and thus impede computational processing — offers a more parsimonious explanation than does the capacity limitation view. Specifically, if early adolescents are more susceptible to interference than older adolescents, they are more likely to find the task of extracting decontextualized representations difficult and to rely on inappropriate representations.

A second argument against the bounded rationality position is based on the framing findings. Instructions to respond from a logical perspective improved performance dramatically. The "logic frame" did not lead the majority of participants to respond normatively and therefore does not rebuke capacity arguments. At the very least, however, the findings indicate that the oft-cited (but rarely clearly defined) capacity boundaries are more flexible than is often implied.

Finally, it was stated earlier that an age-related decrease in the size of the normative/descriptive gap supports the bounded rationality position. However, the observed increases in normatively correct responding should have been accompanied by parallel decreases in fallacious responding. Instead, even the oldest participants responded fallaciously on better than 50% of the problems — heuristic responses were predominant across ages. This finding directly confronts the bounded rationality position and begs for an explanation that theorists from this school of thought will be hard-pressed to generate.

The "inappropriate norms" argument — that individuals are generally rational decision makers, but their decisions are often evaluated against the wrong standards — was also not completely supported. The expectation from this point of view was that the normative/descriptive gap would increase with age. In direct conflict with this hypothesis, the normative/descriptive gap decreased in both frames: If traditional norms are incorrect, why

would norm-consistent responding increase with age and why would responses become more norm-consistent in the logic frame? On the other hand, the finding that fallacies were the predominant responses and did not decline with age seems compatible with the inappropriate norm position. Together, these contradictions, which arose at both the between- and the within-subject levels, imply that both the inappropriate norms and the bounded rationality positions are overly simplistic.

5.2. A two-process explanation

In the two-process theoretic approach adopted here, task representations are key determinants of the information-processing system that assumes predominance in a given situation. Representations, in turn, are determined by various features of the context, including content familiarity, time availability, clues provided by others operating in the same context (e.g., peers), and by characteristics (both relatively stable and transitory) of the individual, including beliefs, perceived needs for accuracy, the ability to inhibit interference, and the competence to generate NCRs.

As in many everyday situations, in most laboratory conditions, participants rely on imprecise gist representations not only because accuracy motivation is low, but also because these representations are unconsciously accepted as accurate reflections of task demands. This conclusion is supported by findings of pervasive overconfidence (e.g., Fischhoff, 1988) and by findings that increasing accuracy motivation (e.g., by informing participants that poor reasoning will be punished) does not reduce certain forms of bias (Klaczynski & Gordon, 1996). Instructions to solve the problems from a ''perfectly logical'' standpoint increased NCRs and decreased biases because, at least for some participants, these instructions cued uncertainty about their normal mode of processing and compelled conscious attempts to process the tasks analytically and to use nondefault representations.

For a variety of reasons, this effort failed for many participants, particularly on the probability and sunk cost problems. For instance, memories activated by task content may have interfered with the abstraction of correct decontextualized representations, interference may have made the correct normative principle difficult to retrieve, a poorly understood or incorrect principle may have been applied, cues for heuristic processing may have overwhelmed attempts to ''switch'' to analytic processing (e.g., by activating powerful judgmental heuristics, such as ''waste not, want not'' on the sunk cost problems), or reasoners equated their ''usual'' representations with ''logical'' representations (see also Reyna et al., in press).

Many of these same explanations could account for the finding that fallacies were pervasive across age groups and were predominant in the usual frame. Given the systematicity with which they appeared, these responses cannot be written off as ''mere'' performance errors; they represent the functioning of a mode of information processing over which individuals have limited control. Because people infrequently monitor their decision making, because negative consequences of heuristic processing are not always immediate or apparent, and because attributions tend to be external when decision outcomes are negative, causal links between decisions and outcomes are made sporadically. As a result, most adolescents

and young adults operate under the assumption that their apparently fallacious reasoning "works." This speculation is supported, to some degree, by indications that the DN, IO, and sunk cost fallacies are well-established by early adolescence and by evidence that fallacious responses in the usual frame rarely changed to NCRs in the logic frame. Specifically, in the logic frame the conditional probabilities of responding normatively on *only one* of the two problems within each JDM task, given fallacious responding in the usual frames were .48, .36, and .10 for the probability, IO, and sunk cost problems, respectively. In other words, when a fallacious response was given in the usual frame, NCRs in the logic frame were uncommon, especially on the sunk cost problems.

The framing effects extend prior research on precision instructions and other manipulations intended to "debias" reasoning (reviewed in Fischhoff, 1988). First, the present enterprise is the first of its type to examine the efficacy of such debiasing manipulations on early and middle adolescent judgments and decisions. Second, as in precision instruction research, but with tasks that have not been examined by debiasing researchers, framing did not eradicate biases and fallacies. Together with precision instruction research, this finding indicates limits on the extent to which "logic framing" can improve normative responding and that such framing is unlikely to influence responses on the types of task studied by Fischhoff (e.g., hindsight bias, overconfidence) and others (see Epstein & Pacini, 1999). Although shifting from inappropriate representations (usual frame) to appropriate representations (logic frame) may be easier on sunk cost, probability judgment, and contrary-to-fact tasks than on hindsight, overconfidence, and conjunction tasks, determining the mechanisms underlying such between-task differences remains an important issue for future research.

5.3. Conclusion and implications for intervention

Two-process theories of decision making are in their infancy and, in many ways, are theoretically underspecified. For instance, heuristic processing, which is thought to occur unconsciously, activates beliefs, social schemata, and easily accessible problem-solving strategies stored in memory. Problem content, internal states, and contextual cues determine which of these types of representation is activated in a particular situation. Subsequent to activation, it is sometimes assumed that these are applied automatically, but it has also been suggested that heuristic strategies appear in consciousness (at least momentarily) and are applied because they "feel" correct (Epstein, 1994; Sloman, 1996). This confusion may be due to the failure of two-process theorists to distinguish between heuristic *processes* (e.g., the unconscious mechanisms that enable peripheral processing of decision-making cues and that activate implicit memory schemata) and heuristic *products* (e.g., judgmental heuristics such as "waste not, want not"). Such distinctions are essential for numerous reasons (Chen & Chaiken, 1999). For instance, when heuristic processing "brings to mind" judgmental heuristics, metacognitively predisposed individuals may operate on, evaluate, and refine these heuristics analytically.

Focusing, for the present, on the *products* of heuristic processing, it can be argued that many heuristics and biases have a cultural basis. DN, the IO fallacy, and the sunk cost fallacy

are repeatedly conveyed to adolescents as appropriate forms of logical argumentation. Agents for the cultural transmission of heuristics include the media (e.g., in life insurance and mutual fund advertisements, anti-drug campaigns), parents (e.g., "If only you'd listened to me . . ."), teachers ("The large university has so many more excellent teachers"), peers ("Look you already spent US$5 on the beer. You may as well drink the whole six-pack. If you don't, you're be throwing your money away."), and other propagators of traditional beliefs (e.g., self-help books promoting such idioms as, "waste not, want not").

It is consistent with psychological tradition, and convenient for those who seek quick solutions for a society plagued by an apparent decision-making crisis among its youth, to assume that optimal decision making involves precision, accuracy, and conscious choice. However, as illustrated in the Introduction, there are conditions under which computational reasoning is not only cognitively wasteful and time-consuming, but is also useless.

Two-process approaches to decision making thus stand in clear disagreement with theories espousing analytic processing as the optimal means by which good decisions are made. Instead, two-process approaches view both cognitive economy and precision as viable decision-making goals, consider "insight" as critical to JDM as is analytic, step-by-step thinking, and argue that the default means by which information is processed is unconscious. In the rapid ebb and flow of everyday life, analytic processing is a luxury we can sometimes afford, although it is required infrequently.

Nonetheless, the prevalence of heuristic responding is not, in itself, an indicator of its adaptive value. The overwhelming frequency with which various heuristics, biases, and reasoning fallacies occur may, in some cases, represent overgeneralizations of automatic responses that are *often*, but not always, useful. Unfortunately, few adults, and still fewer adolescents, understand that ignoring denominators and honoring sunk costs sometimes have maladaptive consequences. The high school senior who ignores ratios of good professors/bad professors — and simply focuses on absolute numbers of good professors — may find him- or herself enrolled in the least intellectually stimulating university of those he or she could have attended. An adolescent girl who continues investing effort to "save" an abusive relationship because the 2 years of her life spent in the relationship will be "wasted" otherwise, is not only putting herself at-risk, but is also ignoring decision options that could produce more satisfactory outcomes.

If this view is correct, interventions should focus not only on facilitating adolescents' understanding of normative reasoning principles, but also on the interdependencies between reasoning and context. Under what conditions should analytic processing be used and normative principles applied? When is heuristic processing more likely to yield adequate decisions than analytic processing? In what types of situations are analytic and heuristic processing likely to produce similar decisions? Only after additional elaborations of two-process theories and further research will it be possible to answer these questions.

References

Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.

Arkes, H. R., & Ayton, P. (1999). The sunk cost and concorde effects: are humans less rational than lower animals? *Psychological Bulletin*, *125*, 591–600.

Bargh, J. A., & Chartland, T. L. (1999). The unbearable automaticity of being. *American Psychologist*, *54*, 462–479.

Baron, J. (1985). *Rationality and intelligence*. Cambridge, MA: Cambridge Univ. Press.

Baron, J. (1988). *Thinking and deciding*. New York: Cambridge Univ. Press.

Brynes, J. P. (1998). *The nature and development of decision making*. Mahwah, NJ: Erlbaum.

Chen, S., & Chaiken, S. (1999). The heuristic–systematic model in its broader context. In: S. Chaiken, & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 73–96). New York: Guilford.

Cohen, L. J. (1981). Can human irrationality be experimentally demonstrated? *Behavioral and Brain Sciences*, *4*, 317–370.

Davidson, D. (1995). The representativeness heuristic and the conjunction fallacy in children's decision making. *Merrill-Palmer Quarterly*, *41*, 328–346.

Denes-Raj, V., & Epstein, S. (1994). Conflict between intuitive and rational processing: when people behave against their better judgment. *Journal of Personality and Social Psychology*, *66*, 819–829.

Epstein, S. (1994). Integration of the cognitive and psychodynamic unconscious. *American Psychologist*, *49*, 709–724.

Epstein, S., Lipson, A., Holstein, C., & Huh, E. (1992). Irrational reactions to negative outcomes: evidence for two conceptual systems. *Journal of Personality and Social Psychology*, *62*, 328–339.

Epstein, S., & Pacini, R. (1999). Some basic issues regarding dual-process theories from the perspective of cognitive–experiential self-theory. In: S. Chaiken, & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 462–482). New York: Guilford.

Evans, J.B.St.T. (1996). Deciding before you think: relevance and reasoning in the selection task. *British Journal of Psychology*, *87*, 223–240.

Evans, J.B.St.T., & Over, D. E. (1996). *Reasoning and rationality*. Hove, UK: Psychology Press.

Fischhoff, B. (1988). Judgment and decision making. In: R. J. Sternberg (Ed.), *The psychology of human thought* (pp. 153–187). Cambridge, England: Cambridge Univ. Press.

Frisch, D. (1993). Reasons for framing effects. *Organization Behavior and Human Decision Processes*, *54*, 399–429.

Gigerenzer, G. (1996). On narrow norms and vague heuristics: a reply to Kahneman and Tversky. *Psychological Review*, *103*, 592–596.

Horn, J. L. (1970). Organization of data on life-span development of human abilities. In: L. R. Goulet, & P. B. Baltes (Eds.), *Life-span developmental psychology: research and theory* (pp. 424–466). New York: Academic Press.

Jacobs, J. E., & Potenza, M. (1991). The use of judgment heuristics to make social and object decisions: a developmental perspective. *Child Development*, *62*, 166–178.

Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, *39*, 581–582.

Klaczynski, P. A. (2000). Motivated scientific reasoning biases, epistemological beliefs, and theory polarization: a two-process approach to adolescent cognition. *Child Development*, *71*, 1347–1366.

Klaczynski, P. A., & Gordon, D. H. (1996). Self-serving influences on adolescents' evaluations of belief-relevant evidence. *Journal of Experimental Child Psychology*, *62*, 317–339.

Klaczynski, P. A., & Narasimham, G. (1998). Problem representations as mediators of adolescent deductive reasoning. *Developmental Psychology*, *34*, 865–881.

Kreitler, S., & Kreitler, H. (1986). Development of probability thinking in children 5 to 12 years old. *Cognitive Development*, *1*, 365–390.

Margolis, H. (1987). *Patterns, thinking, and cognition*. Chicago: University of Chicago Press.

Oaksford, M., & Chater, N. (1993). Reasoning theories and bounded rationality. In: K. Manktelow, & D. Over (Eds.), *Rationality: psychological and philosophical perspectives* (pp. 31–60). London: Routledge.

Reber, A. S., Walkenfeld, F. F., & Hernstadt, R. (1991). Implicit and explicit learning: individual differences and IQ. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 888–896.

Reyna, V. F., & Brainerd, C. J. (1995). Fuzzy-trace theory: an interim synthesis. *Learning and Individual Differences*, *7*, 1–75.

Reyna, V. F., & Ellis, S. C. (1994). Fuzzy-trace theory and framing effects in children's risky decision making. *Psychological Science*, *5*, 275–279.

Reyna, V. F., Lloyd, F. J., & Brainerd, C. J. (in press). Memory, development, and rationality: an integrative theory of judgment and decision making. In: S. Schneider & J. Shanteau (Eds.), *Emerging perspectives on decision research*. Cambridge Univ. Press.

Siegler, R. S. (1996). *Emerging minds: the process of change in children's thinking*. New York: Oxford Univ. Press.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*, 129–138.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*, 3–22.

Stanovich, K. E. (1999). *Who is rational? Studies of individual differences in reasoning*. Mahwah, NJ: Erlbaum.

Stanovich, K. E., & West, R. F. (1997). Reasoning independently of prior belief and individual differences in actively open-minded thinking. *Journal of Educational Psychology*, *89*, 342–357.

Stanovich, K. E., & West, R. F. (1998). Individual differences in rational thought. *Journal of Experimental Psychology: General*, *127*, 161–188.

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: implications for the rationality debate? *Behavioral and Brain Sciences*, *23*, 645–726.

Stein, E. (1996). *Without good reason: the rationality debate in philosophy and cognitive psychology*. Oxford: Oxford Univ. Press.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*, 453–458.