

Refusing to budge: a confirmatory bias in decision making?

Lea-Rachel D. Kosnik

Received: 15 January 2007 / Accepted: 16 October 2007
© Fondazione Rosselli 2007

Abstract Confirmatory bias, defined as the tendency to misinterpret new pieces of evidence as confirming previously held hypotheses, can lead to implacable, even incorrect decision making. It is one of the biases, along with anchoring, framing, and other judgment heuristic errors, that may lead to non-optimal behavior. This paper tests for the existence of confirmatory bias behavior in a uniquely economic setting (tax policy) and in a context relatively lacking in ambiguity. It also tests whether the confirmatory bias phenomenon can be prevalent enough to affect aggregate outcomes, a characteristic important in economic models in particular. The results indicate not only that confirmatory bias exists, but that the confirmatory bias effect may be stronger for evidence relating to losses than for comparable evidence relating to gains.

Keywords Confirmatory bias · Rationality · Judgment · Heuristics · Tax policy

1 Introduction

The relevance, and prevalence, of particular judgment biases in decision making continues to be a subject of debate in a variety of disciplines. Exploring the applicability of these biases in disparate fields and under varying circumstances is an active area of modern research. This paper contributes to the discussion by focusing on one judgment bias in particular, confirmatory bias in evidence

L.-R. D. Kosnik (✉)
Department of Economics,
University of Missouri—St Louis,
8001 Natural Bridge Road,
St Louis, MO 63121, USA
e-mail: kosnikl@umsl.edu

evaluation and decision making, and investigating it in a uniquely economic context.

Confirmatory bias occurs when, unlike rational Bayesian actors, people misinterpret newly presented information as supporting previously held beliefs, even when there is no clear evidence for such an interpretation.¹ A classic example from the psychology literature involves evaluation rates of pupils by teachers. Once a teacher has formed an opinion regarding the skill of a pupil, the argument goes, further pieces of evidence on the abilities of that pupil (say, an essay) can be used to support both a positive, or a negative, assessment of that very same student, depending on the initial beliefs of the instructor. Investors, confident in their evaluations of particular financial theories (the efficiency of markets as opposed to active market involvement) may interpret new evidence (a revaluation of the Chinese Yuan) as supportive of each of their independent views. The arrest and prosecution of Martha Stewart in the United States, for example, may be perceived as evidence, on the one hand, of a crackdown on white-collar crime, or, as evidence, depending on your prior beliefs, of the relative paucity of white-collar crime arrest rates. As succinctly stated by Rabin and Schrag (1999, p. 38), “A person suffers from confirmatory bias if he tends to misinterpret ambiguous evidence as confirming his current hypotheses about the world.”

If confirmatory bias exists and exerts consistent, significant pressure on the evaluations of a wide range of economic actors (from teachers, to investors, to interpreters of current events), its prevalence can lead to poor decision making behavior. Teachers may penalize perfectly able students, investors may make poor financial decisions and citizens may vote without valid comprehension of an issue. Rather than rational Bayesian updating in an uncertain world, economic actors actually misread or disregard newly acquired pieces of information primarily on the basis of its harmony (or disharmony) with previously held opinions.

It is sensible, of course, to use newly available evidence to update one's prior beliefs, and it is also sensible to avoid letting a single piece of evidence disproportionately sway previously well-thought-out hypotheses, but confirmatory bias is about the direction of causation mistakenly running the other way. It is the error of using a priori beliefs to then subjectively evaluate new pieces of evidence. This distinction is fundamentally important, and often misunderstood in discussions of the confirmatory bias phenomenon. Economists and other utility maximizing theorists employ the basic assumption that rational actors evaluate all new pieces of information in an unbiased, nonsubjective, cool and level-headed way. Confirmatory bias is, in essence, when an actor instead evaluates new pieces of information from the world around him in a subjective, misguided fashion, based on prior opinion. While it is, at heart, an informational processing problem, confirmatory bias cannot be dismissed as a transaction cost problem in light of its nonsymmetric nature; confirmatory evidence is kept and processed, regardless of the difficulties in digesting and translating such information into a coherent framework, while

¹ Confirmatory bias can also refer, more broadly, to control strategies where people seek (and not just interpret) biased information when forming initial hypotheses, but that separate branch of the literature is not the focus of this study.

disconfirmatory evidence, however simplistic, is summarily misread and often disregarded.²

Psychological evidence for the existence of the confirmatory bias phenomenon is abundant (Bruner and Potter 1964; Lord et al. 1979; Plous 1991; Darley and Gross 1983; Haverkamp 1993; Dougherty et al. 1994; Strohmer and Shivy 1994; Snyder and Swann 1978; Swann and Snyder 1980; Miller et al. 1993; Nickerson 1998), but aside from Rabin and Schrag's (1999) theoretical modeling of the bias in an economic methodological framework, there has been little empirical or experimental testing of the concept geared towards an economic setting.^{3,4} This paper is an attempt to address that oversight.

Further, this paper seeks to test the "seriousness" of the confirmatory bias concept; in other words, this paper is also a test of the conceptual boundaries for assessment of confirmatory bias behavior. Hitherto, evidence in support of confirmatory bias was sought through an actor's treatment of uncertain or "ambiguous" evidence (Keren 1987, 1988; Griffin and Tversky 1992; Rabin 1998; Rabin and Schrag 1999; Lord et al. 1979; Dunning et al. 1989). For example, previous studies of confirmatory bias have used as evidence newspaper articles, academic essays, or video clips, all of which can not be analyzed subject to precise mathematical calculations. The rigorous Rabin and Schrag (1999) model, however, shows that confirmatory bias may lead actors to make outright "wrong" decisions. Liberally, this result could be taken to imply that an agent suffering from confirmatory bias may apply this judgment error to relatively unambiguous data as well as ambiguous data (albeit, perhaps less dramatically, for as Rabin and Schrag note, this "wrongness" is only exhibited in the extreme).

In practical terms, given less subjective evidence on a pupil's performance (multiple choice exams with definitively correct answers, as opposed to an essay exam), a teacher may still misread this less ambiguous data if it fails to conform to prior beliefs. Brokers who lose substantial sums of money for their clients, year after

² It may be useful to think of the problem as loosely akin to a "selectivity bias" or "added-variable bias" in information processing: the idea is that sometimes people misread evidence relevant to a particular hypothesis and either misread this evidence by refusing to admit its contrary nature, thereby disregarding it as useless (selectivity bias), or people misread the evidence by processing it as supportive of previously held views even when it is not (added-variable bias). Either way, such data manipulation leads to biased estimates (both in support of the a priori hypothesis) and poor decision-making behavior.

³ Overconfidence, a behavioral phenomenon related to confirmatory bias, has been subject to some scrutiny in the economics literature (see, for example, Ganguly et al. 2000; Madsen 1994; Camerer 1997; Camerer and Lovo 1999; Eales et al. 1990; Dubra 2004; Hvide 2002; Golec and Tamarkin 1995; Barber and Odean 2001), but similar attention has not been paid to the confirmatory bias effect itself.

⁴ While it appears that direct tests of the confirmatory bias phenomenon have yet to be published in the economics literature, it is possible to reevaluate previously published work in light of the confirmatory bias effect. For example, in an historical article documenting the pervasive extent of homestead failures in the Upper Great Plains between 1890 and 1925, Libecap and Hansen (2000) come to the conclusion that the widespread homesteading failures, which resulted in significant monetary and nonmonetary losses, were made because: "...the previous wet period and the strong claims of dryfarming experts led homesteaders to discount observations of dry weather and to place more weight on past opinions about the ability of the region to withstand droughts" (p. 28). This conclusion, while reasonable, begs the question as to *why* such a misreading of the available information took place. A strong case could be made that this is evidence of confirmatory bias behavior. It is quite likely that other papers in a similar vein could also be reinterpreted—and better understood—under the confirmatory bias lens.

year, may still find support from their investors out of “loyalty”, a “belief in the underdog”, or some other reasoning that expressly goes against the not very ambiguous evidence of consistent financial losses.⁵ Admittedly, what constitutes ambiguous versus unambiguous evidence is a bit subject to ambiguity itself (were the multiple choice questions written poorly? how much money does have to be lost before it is considered “substantial”?), but this is where an economic framework for testing the existence of confirmatory bias comes in particularly handy. By using tax policy as the subject of the current experiment, measuring the effects of particular evidence to an assumed utility maximizing actor are less susceptible to ambiguous interpretation (a \$1,000 tax increase is, regardless of any indirect social benefits or personal feelings for tax policies in general, notably utility decreasing). One hypothesis of this paper, then, is that even with much less subjective evidence, some actors will still be prone to the confirmatory bias phenomenon. This is in stark contrast to the predictions of traditional utility maximizing theory.

Finally, even if we discover that confirmatory bias exists in economic settings and with less ambiguous evidence, the relevance of such a finding would only be worthwhile (in particular, to economists) if it affected aggregate welfare outcomes. If a few actors make a few extreme mistakes, does it matter? Our final test, therefore, is not simply on the existence of confirmatory bias, but also on its prevalence.

The experimental data for this paper come from a 2004 survey on tax policy in the United States given to students in two sections of an introductory economics class at Montana State University, Bozeman. Tax policy (in particular, of the second Bush Administration) was chosen as the contextual subject for the survey because it satisfies three important constraints: (1) it is an economic policy issue, (2) taxes tend to elicit strong policy opinions from at least a small subset of the population and confirmatory bias is hypothesized to exist only for decisions involving a strong degree of prior beliefs, and (3) tax policy was believed to be a particularly credible survey instrument in light of the active tax-cutting agenda of the current Bush Administration.

The survey was modeled on previous designs targeted to test the existence of possible confirmatory bias behavior (Lord et al. 1979, 1984; Plous 1991). Details of the survey are outlined in Sect. 3 of this paper. The results indicate affirmative evidence for the existence of confirmatory bias behavior within the context of United States tax policy, for data that is notably less ambiguous than that used in previous research, and for a significant enough number of actors to affect aggregate outcomes. Further, and unexpectedly, the results indicate that confirmatory bias may be stronger for evidence related to a tax increase (which reduces subjects’ utility),

⁵ In another context, Ausubel (1991) finds that, despite consistent, unambiguous evidence of credit card debt (monthly statements), many credit card holders (47%) insist that they “nearly always” pay their balances in full (with an additional 26% claiming that they “sometimes pay” their balances in full). The fact is that three-quarters of all active credit card accounts at major banks average over \$1,000 in balances at any moment in time. This belief by debt holders in their ability to pay off their credit cards every month and not incur continuous, often exorbitant finance charges, flies in the face of their (unambiguous) monthly available evidence. It can also lead them to incur further, perhaps nonoptimal, extra credit card debt.

than for evidence on a tax decrease (which increases subjects' utility). These results tentatively suggest not only that confirmatory bias exists, but that it may be stronger for losses (prospect theory, for example, assumes that losses are treated differently than gains) than for gains.

Section 2 presents a brief review of confirmatory bias as it has been examined in the psychology and economic literatures, Sect. 3 goes over the design of the survey instrument used in this research, Sect. 4 discusses the empirical results, and Sect. 5 sums up the conclusions.

2 Literature and methodological review

Cognitive studies on confirmatory bias have identified two main manifestations of the behavior: (1) confirmatory bias as a control strategy, and (2) confirmatory bias in the interpretation of new information. Confirmatory bias as a control strategy is most relevant to hypothesis formation and testing, when an actor seeks hypothesis-related information that is primarily positive or confirmatory, often in an asymmetrically diagnostic way. Seminal studies include Wason (1960), Tweney et al. (1980), Klayman and Ha (1987) and Trope and Mackie (1987). While an interesting branch of the literature, this paper focuses on the second main category of confirmatory bias behavior, that of interpretation of already selected information.

Confirmatory bias and the practice of assimilating new pieces of *given* information primarily in keeping with one's prior beliefs was documented in a seminal study by Lord et al. (1979). After isolating a group of subjects who exhibited strong prior beliefs on the deterrent efficacy of capital punishment (either in support of, or against), these subjects were then presented with an identical body of relevant evidence (two academic studies, one seemingly confirming and the other disconfirming the deterrent efficacy of the death penalty) and were asked to rate the quality and persuasiveness of the two studies.⁶ Both proponents and opponents of capital punishment believed that the results and procedures that confirmed their own beliefs were of better quality, relevance, and persuasiveness than the studies that did not. Subsequently, the subjects reported shifts in their beliefs regarding capital punishment that resulted in greater polarization of attitude than at the start of the study. As the authors note, "If people of opposing views can each find support for those views in the same body of evidence, it is small wonder that social science research...will frequently fuel rather than calm the fires of debate." Confirmatory bias, in other words, is self-propelling and can lead to attitude polarization on particularly contentious or emotional issues. Other studies illustrating the prevalence of confirmatory bias and attitude polarization in a range of settings include Swann and Snyder (1980) and Darley and Gross (1983) with application to student evaluations, Dougherty et al. (1994) with application to evaluating job candidates,

⁶ Keep in mind that the subjects are being asked to evaluate only the new evidence. They are not (initially) being asked to reevaluate their overall opinion of capital punishment. Once the new evidence is assimilated, subjects may very well then use it to update their beliefs regarding capital punishment, but it is incorrect according to the traditional neo-classical model to use prior beliefs to subjectively analyze newly available evidence.

Strohmer and Shivy (1994) and Haverkamp (1993) with application to counselor/client evaluations, and Miller et al. (1993) with applications to capital punishment and affirmative action.

Recently, Rabin and Schrag (1999) modeled the confirmatory bias phenomenon in a mathematical framework and discovered some intriguing implications. Given two states of the world, $x \in \{A, B\}$, where A and B are exhaustive and mutually exclusive hypotheses concerning a particular issue, an agent forms beliefs on the likelihood of each state of the world from signals that he receives over time. In every period $t \in \{1, 2, 3, \dots\}$ the agent receives a signal, $s_t \in \{a, b\}$ that is correlated with the true state of the world. These signals are independent and identically distributed, with $\text{prob}(s_t = a|A) = \text{prob}(s_t = b|B) = \theta$, for some $\theta \in (0.5, 1)$. After observing each signal, the agent updates his belief (using Bayes' Rule) about the relative likelihood of $x = A$ and $x = B$.

To model confirmatory bias, Rabin and Schrag suppose that the agent may misread (with probability $q > 0$; the greater is q the more extreme the bias) signals that conflict with his current belief about which state of the world is more likely. If he believes that state A is more likely than state B , he may misread a conflicting signal $s_t = b$ believing instead that he observes $s_t = a$.⁷ The presence of confirmatory bias means that the agent's perceived signals are neither independently nor identically distributed. In general, when $q > 0$, an agent's beliefs in terms of a relative likelihood ratio are biased; if he believes in state A due to misperceived signals in support of A , he is overconfident.⁸

Confirmatory bias, it appears, leads to another well documented behavioral phenomenon, overconfidence: given any probabilistic assessment by a player that a certain hypothesis is true, appropriate beliefs would in fact deem it less likely to be true. Studies have shown, for example, that most people (in one study, 90%) believe they are above average drivers (Svenson 1981).⁹ Another study found this to be true despite learning and accumulated evidence.¹⁰ Studies providing evidence for the overconfidence effect within economic contexts in particular include Ganguly et al. (2000) and Barber and Odean (2001), who provide evidence for an overconfidence effect in asset market evaluations and stock trading, Madsen (1994) who finds overconfidence effects in manufacturing data, Camerer and Lovallo (1999) who provide experimental evidence that overconfidence leads to excessive entry decisions in the marketplace, Eales et al. (1990) who find that overconfidence biases Illinois farmers and grain merchandisers evaluations of option prices, and

⁷ This structure conforms to the "added-variable" type of bias in informational processing mentioned in the introduction.

⁸ See Rabin and Schrag (1999) for a fuller exposition of the model, including situations of possible underconfidence.

⁹ Other applications have been made to management consulting (Watson et al. 1998), medical practitioners (Arkes et al. 1981) forensic evaluation (Borum et al. 1993) and prediction (Kahneman and Tversky 1973). Cutting closer to home, Cross (1977) has found that 94% of college professors (drawn from a wide distribution of disciplines) say they do above average work. An informal survey of economic professors at the 2004 AEA conference found similar ability overconfidence levels, as well as a professed readiness to bet real money on those beliefs.

¹⁰ Preston and Harris (1965) found significant overconfidence regarding driver ability from respondents surveyed in a hospital just after a bad car accident.

Golec and Tamarkin (1995) who find evidence for the overconfidence effect influencing betting behavior. While the purpose of this paper is to test for evidence of confirmatory bias, not its subsequent manifestation of overconfidence, these studies all point to the economic relevance of this issue for behavioral decision making.

Even more tantalizing than the overconfidence results, Rabin and Schrag (1999) show that an agent may suffer not merely from overconfidence, but also from “wrongness,” when an agent’s confirmatory bias is severe enough. In other words, agents suffering from confirmatory bias (in evidence evaluation) may not merely make hubristic, suboptimal, but generally on-target mistakes, they may make outright incorrect, incontrovertibly utility reducing decisions. This is the theoretical motivation for testing, in this paper, the hypothesis that confirmatory bias behavior exists even when assimilating (relatively) unambiguous data. If confirmatory bias is severe enough, actors may make outright “wrong” decisions, despite the evidence. In the extreme, Rabin and Schrag show that an agent exhibiting confirmatory bias may come to believe with near certainty in a false hypothesis despite receiving an infinite amount of information. In contrast, a Bayesian actor would, after an infinite amount of information, believe with near certainty in the correct hypothesis. No matter how much money is lost, some investors may stick with a poor financial analyst despite the reams of contrary evidence.¹¹ We would assume that the number of actors who do this would not be the majority (otherwise, they would be “money pumps” or “suckers”), but they may still be enough to affect overall outcomes.

Speculation as to why confirmatory bias might exist, as opposed to if it exists, has been more limited. Some studies have justified confirmatory bias-type behavior through self-assessment benefit-cost calculations (Babcock and Loewenstein 1997; Akerlof 1989; Taylor and Brown 1988). According to Dunning et al. (1989, p. 1082), “...the world is at times an uncaring place, and to maintain a happy and productive life, one must color one’s judgments with optimistic views that are less than completely grounded in reality.” In other words, the benefits of confirmatory bias behavior in positively affecting one’s outlook on the world may outweigh any costs in judgment error. It is common in the psychology literature to come across arguments that imply people are happier, less depressed, more confident, even more successful, if they can believe they are right most of the time and if they can assimilate evidence from the world around them as supporting their personal beliefs.¹² Similarly, Dardenne and Leyens (1995) make the argument that confirmatory bias is a social skill, adaptively useful in identifying similarities with others in certain social situations. Kahneman and Tversky (1973) and

¹¹ This implies, notably, that learning and experience will not automatically correct an agent’s misperceived beliefs. As Baumann et al. (1991) note, agents must first perceive a need for change in order for learning and experience to effect change, and confirmatory bias implies that actors do not perceive anything wrong with their biasedly confident decision making.

¹² In a similar vein, overconfidence may be used not just to fool ourselves into being happier, more productive individuals, overconfidence may sometimes even be used to fool others. As an example, an overconfidence in one’s defensive abilities may lead to challenges of a much more powerful enemy. Such hubris, rather than leading to defeat, may actually cause the rival to question herself and ultimately back down without a fight, thereby handing the victory to the less able, but overconfident, aggressor. The author thanks Rob Fleck for pointing out this natural selection application of the overconfidence effect.

Kahneman et al. (1982), meanwhile, explain many purported judgment errors similar to confirmatory bias as pragmatic behavior in light of humans' limited cognitive abilities. Other research (Nisbett and Ross 1980) has also led to the conclusion that humans are woefully inept information processors and that this may be a reason why biases such as confirmatory bias persist. A greater understanding of the origins of confirmatory bias and why it might develop would be a useful area for future research; in the meantime, the immediate relevance of this study is in its suggestion that confirmatory bias does exist, that it has significant effects, and also that its influence may be greater in certain directional contexts than in others.

3 Survey design

A total of 284 participants took part in a confidential survey on tax policy in the United States. Limited demographic data on the subject pool was collected (Table 1), and as can be seen from the summary statistics, the subject pool was a relatively homogenous lot consisting primarily of young, low income, single, white, US citizens. The questionnaires were administered in a group setting and, for motivation, extra credit points were awarded as students turned their completed (and anonymous) questionnaires in to the administrator at the end of the survey periods.¹³

The survey was designed and the questions were chosen as much as possible to mimic previous surveys in the literature used to test for the existence of confirmatory bias behavior (Lord et al. 1979, 1984; Plous 1991). The survey began with a list of four questions (Appendix 1), all culled or adapted from national opinion polls by Gallup, CNN, or USA Today, that asked respondents to state their advocacy position regarding current tax policy in the United States. The answers were rated on a five-point scale (with (5) indicating strong, positive support for current tax policy, and (1) strong, negative support) and averaged to find an initial predisposition to current Bush Administration tax policy.¹⁴ The (normally distributed) results indicate that this subsection of the population had slightly pro-Bush (>3) initial proclivities.¹⁵

In Part II of the survey, the subjects were asked to read about a new tax policy proposition of the Bush Administration.¹⁶ This was the “evidence” relating to current tax policy in the United States. Half of the subjects were presented with a

¹³ The data from the two survey periods were combined to form a single comprehensive data set for use in all future data analysis, after first confirming through statistical analysis (an insignificant dummy variable proxying for the two sections) that the data were formally consistent.

¹⁴ In order to reduce survey bias, the order of the answer choices among the questions and among the respondents was mixed; half the subjects received answer choices moving from “greatly helped...” or “greatly weaken...” vertically down the page to “greatly hurt...” or “greatly strengthen...” while the other half of the respondents received surveys with the answer choices moving in the other direction.

¹⁵ If we can assume that this subject pool is representative of the entire Montana State University student body, then this revealed proclivity is not surprising; a majority of MSU students are Montana natives, and Montana historically has been a Republican-leaning state.

¹⁶ The separate sections of the survey were clearly delineated in bold as “Part I,” “Part II,” etc., and presented on separate pages of the survey booklet in order to reduce, as much as possible, potential order effects between the questions. This method was suggested in Dillman (2000).

Table 1 Summary statistics

Frequency distributions (%)						<i>n</i> = 276
Variable	1	2	3	4	5	
<i>Demographic Vars.</i>						
Age	13.87	30.66	22.63	12.04	20.8	
Gender	41.24	58.76				
Marital status	93.43	5.47	0.73	0.36		
Race	94.16	2.19	0.36	1.09	2.19	
Annual income	92.7	5.47	1.82			
Parents' education	0.73	22.63	45.99	30.66		
US citizen	97.08	2.92				
Filed US taxes	83.15	16.85				
First year filed US taxes	23.58	12.74	21.70	19.34	22.64	
IRS audit	2.94	97.06				
Variable	Mean	Standard deviation	Minimum	Maximum		
<i>Survey Vars.</i>						
Part I: Initial predisposition						
Question 1	3.26	1.04	1	5		
Question 2	3.07	1.08	1	5		
Question 3	3.40	1.08	1	5		
Question 4	2.58	0.90	1	5		
Part II: Evidence assimilation						
Question 5	1	0	1	1		
Question 6	0.87	0.34	0	1		
Question 7	0.98	0.15	0	1		
Question 8	3.39	1.46	1	5		
Part III: Attitude polarization						
Question 9	-0.15	0.88	-2	2		
Question 10	-0.00	1.10	-2	2		
Question 11	-0.02	0.95	-2	2		

Age: 1 = 18, 2 = 19, 3 = 20, 4 = 21, 5 = ≥ 22

Gender: 1, female; 2, male.

Marital status: 1, single/never married; 2, married; 3, separated/divorced; 4, widowed

Race: 1, white; 2, black or African American; 3, American Indian and Alaska native; 4, Asian; 5, some other race

Annual income: 1, \$20,000 or less; 2, \$20,001–\$40,000; 3, \$40,001–\$75,000; 4, \$75,001 or more

Parent's highest education level: 1, less than 12th grade; 2, high school with diploma; 3, college/vocational training; with degree; 4, post-college graduate or professional training with degree

Are you a US Citizen: 1, yes; 2, no

Have you ever filed US federal income taxes before: 1, yes; 2, no

If filed, in what year did you first begin filing tax returns: 1, ≤ 1998 ; 2, 1999; 3, 2000; 4, 2001; 5, ≥ 2002

Have you ever been audited by the IRS before: 1, yes; 2, no

policy scenario (Survey A) indicating a tax reduction, in the form of an increase of the Hope Scholarship Credit for college students. The other half of the subjects were given a policy scenario (Survey B) describing a tax increase, this time due to a reduction of the Hope Scholarship Credit for college students. The description of the credit in both scenarios was brief (only two paragraphs), succinct, and to the point. It ended with the sentence: “If you are a current college student, therefore, you will unambiguously benefit (Survey A)/lose (Survey B) from this proposal.” Such a straightforward presentation of the evidence was intended to make it difficult for the subjects to fail to recognize the utility maximizing implications of the proposal.

To assure comprehension, subjects were asked to complete a series of multiple-choice questions on the details of the presented policy proposal (Appendix 1, Q5–Q7). Their answers were scored with a 1 for a correct answer, and 0 for any incorrect choice. Subjects overall exhibited a high degree of motivation, with only 8 out of 284 surveys (2.8%) receiving a combined comprehension score of less than 67% (meaning, more than one out of the three questions wrong). About 84.9% of all subjects received a perfect score of 100%. Those subjects that did score less than a 67% were dropped from the final data analysis.

At the end of the policy analysis, subjects were asked to rate their support or dissent for the new policy proposal, again on a five-point scale (Appendix 1, Q8). To be clear, this question is asking how students assimilated the new policy evidence only. It is not asking about support or opposition for the notion of taxes in general, nor is it asking for support or opposition for the Bush Administration in general. The question specifically focuses on the “latest” evidence, and not on ultimate support for or against the Bush Administration. In other words, it is logically possible for people with strong anti-Bush policy feelings to still respond with strong support in favor of the tax cut. Indeed, from a traditional utility maximizing perspective, that is exactly what they would be predicted to do. If a subject instead conflates Q8 with other issues, particularly with opinions regarding the Bush Administration in general, that could constitute evidence of the directional causation error of confirmatory bias behavior.

Finally, to test not merely for biased assimilation of the presented evidence, but also for subsequent attitude polarization, the subjects answered a series of questions designed to assess attitude change towards tax policy in the Bush Administration (Part III, Appendix 1).¹⁷ The answers to these questions were rated on a scale from –2 to +2 (a negative (positive) score indicating decreased (increased) support for current tax policy, and a score of zero indicating no change in attitude due to the presented evidence) and then averaged to form a composite unweighted index of overall change in advocacy position.

Results of the survey are presented below.

4 Results and discussion

If agents were fully rational in the neo-classical sense and unaffected by confirmatory bias, then traditional utility maximization theory would predict that

¹⁷ This is the point at which it is correct to use the new evidence to update prior beliefs, but only, of course, if that new evidence was first analyzed objectively and independently.




Fig. 1 Correlation between initial predisposition and evidence assimilation

the subjects of this survey would have strongly supported the tax cut in Survey A (a numerical index of 5) and opposed the tax increase in Survey B (a numerical index of 1), regardless of their initial predispositions to the Bush Administration's tax policy plans in general. This null hypothesis stems from the traditional assumption by economists that "individuals vote in their own best interest" (Akerlof 1989, p. 10), (Peltzman 1980).^{18,19} Figure 1 reports the means and first standard deviation of support or opposition to the tax policy evidence presented in Surveys A and B, given initial predisposition to Bush Administration tax policy. A correlational trend (especially strong in Survey B) relating initial predisposition and subsequent assimilation of the survey evidence is clearly visible.²⁰

To test more specifically for evidence of confirmatory bias behavior, we first note that the current definition of confirmatory bias in evidence evaluation assumes that agents only exhibit this bias with decisions for which they have a strong initial proclivity.²¹ If you are vehemently against (or for) capital punishment, for example, you will view most evidence you encounter subsequent to forming this particular

¹⁸ Even models that allow for an altruism component of voting behavior (Nelson 1994), never suggest that this component is strong enough to overwhelm the objective of maximization of real income, although, altruism may dampen it.

¹⁹ An alternative interpretation of subjects' response may be related to ideology's commitment enhancing role in settings of uncertainty (Downs 1957), but that explanation lacks teeth as there is little uncertainty in this particular setting, which is the point. The relatively unambiguous survey evidence makes it directly clear to subjects how these policy proposals will affect their ultimate tax returns. Put another way, there are two types of uncertainty which may affect actors' behavior: uncertainty with the datum and uncertainty of preferences. As in most of economics, this paper assumes away uncertainty of preferences. That leaves only uncertainty in the data which the survey evidence was explicitly designed to avoid.

²⁰ The author thanks an anonymous reviewer for pointing out that the more linear relationship observed in the Survey A graph, as opposed to the Survey B graph, may be due to a ceiling effect.

²¹ This is not necessarily true for confirmatory bias stemming from control strategies, such as with Wason's 2-4-6 and other selection tasks, where strongly held beliefs are not required for biased hypothesis updating.

Table 2 Evidence assimilation: tabular results on restricted sample

	Survey A—tax cut ($n = 140$)			Survey B—tax increase ($n = 136$)		
	Evidence assimilation			Evidence assimilation		
Initial predisposition	In favor of current tax proposal	Neutral regarding current tax proposal	Against current tax proposal	In favor of current tax proposal	Neutral regarding current tax proposal	Against current tax proposal
Strong support for Bush Administration tax policy	$\frac{57}{63} = (91\%)$	$\frac{4}{63} = (6\%)$	$\frac{2}{63} = (3\%)$	$\frac{17}{48} = (35\%)$	$\frac{11}{48} = (23\%)$	$\frac{20}{48} = (42\%)$
Strong opposition to Bush Administration tax policy	$\frac{13}{21} = (62\%)$	$\frac{7}{21} = (33\%)$	$\frac{1}{21} = (5\%)$	$\frac{2}{25} = (8\%)$	$\frac{5}{25} = (20\%)$	$\frac{18}{25} = (72\%)$
	Chi ² (2) = 10.41; $Pr = 0.005$			Chi ² (2) = 7.72; $Pr = 0.021$		

Strong **support** for Bush Administration tax policy = initial predisposition score ≥ 4

Strong **opposition** to Bush Administration tax policy = initial predisposition score ≤ 2

In **favor** of current tax proposal = evidence assimilation score ≥ 4

Neutral regarding current tax proposal = evidence assimilation score = 3

Against current tax proposal = evidence assimilation score ≤ 2

opinion as supportive of your initial point of view. (In other words, if you strongly believe that $x = A$, you will perceive most subsequent signals as $s_t = a$.)

For this reason, most of the published empirical literature testing for evidence of confirmatory bias first isolates those subjects indicating strong initial beliefs regarding the contextual subject of the survey. In our experiment, this implies concentrating on the subjects whose composite initial predisposition score is 2 or less, or 4 or more. Table 2 reports the frequencies between initial predisposition (rows), and subsequent evaluation of the survey evidence (columns), for those subjects with strong initial beliefs regarding Bush Administration tax policy.²² The Chi-square tests for both surveys find that the differences between frequencies are significant. Note that *Favor* is the neoclassical “rational” response for Survey A, and *Against* is the neoclassical “rational” response to the evidence presented in Survey B. For both surveys, if confirmatory bias exists, evidence for it should correspond to the cells {Support, Favor}, {Support, Neutral} or {Opposition, Against}, {Opposition, Neutral}. This is because if you initially support Bush Administration tax policy (i.e. you strongly believe that $x = A$, where $A = \text{Support}$), your evidence assimilation will correspond in a similarly supportive way (i.e. $s_t = a$, where $a = \text{Favor}$, or at the very least, Neutral). Symmetrically, if your initial predisposition is in opposition to Bush Administration tax policy (i.e. $x = B$, where

²² Note that 84 observations in Survey A and 73 observations in Survey B, or 60 and 54% of the total sample sizes respectively, fit into the two tails of the normal distributions of responses from both surveys.

Table 3 Attitude change—tabular results on restricted sample

	Survey A—tax cut (<i>n</i> = 140)			Survey B—tax increase (<i>n</i> = 136)		
	Attitude change			Attitude change		
Initial predisposition	Positive attitude change	Zero attitude change	Negative attitude change	Positive attitude change	Zero attitude change	Negative attitude change
Strong support for Bush Administration tax policy	$\frac{41}{63} = (65\%)$	$\frac{18}{63} = (29\%)$	$\frac{4}{63} = (6\%)$	$\frac{7}{48} = (15\%)$	$\frac{25}{48} = (52\%)$	$\frac{16}{48} = (33\%)$
Strong opposition to Bush Administration tax policy	$\frac{6}{21} = (29\%)$	$\frac{9}{21} = (43\%)$	$\frac{6}{21} = (29\%)$	$\frac{1}{25} = (4\%)$	$\frac{6}{25} = (24\%)$	$\frac{18}{25} = (72\%)$
	Chi ² (2) = 11.29; <i>Pr</i> = 0.004			Chi ² (2) = 10.01; <i>Pr</i> = 0.007		

Strong **support** for Bush Administration tax policy = initial predisposition score ≥ 4

Strong **opposition** to Bush Administration tax policy = initial predisposition score ≤ 2

Positive attitude change = composite attitude change score > 0

Zero attitude change = composite attitude change score = 0

Negative attitude change = composite attitude change score < 0

B = Opposition), your assimilation of the new evidence will also be in a negative fashion (i.e. $s_t = b$, where b = Against, or at the very least, Neutral).²³

Looking at Survey A, those with an initial Support predisposition in support of Bush Administration tax policy, overwhelmingly assimilated the subsequent tax policy evidence in a similarly supportive manner. This was also the “rational” response, however, so we can not discern whether the responses were based on a calculation of utility maximization or confirmatory bias. Looking at the bottom row of Survey A, however, the Opposition row, those whose initial predisposition was in opposition to Bush Administration tax policy, showed a significantly smaller tendency to pick the “rational” response. Instead, the subjects placed nearly 40% of their assimilation responses into the opposing or neutral category, thereby avoiding the “rational” response in favor of maintaining their initial predisposition.²⁴

The results from Survey B show the same sorts of tendencies. Those who initially opposed the Bush Administration’s tax agenda, subsequently assimilated the new evidence in a similarly oppositional way, while those that initially supported the Bush Administration’s tax policy (the Support row), rationally opposed the tax

²³ Another way of making the same point is that {Support, Favor} and {Opposition, Against} correspond to the added-variable interpretation of confirmatory bias behavior, while {Support, Neutral} and {Opposition, Neutral} relate to the selectivity bias interpretation of confirmatory bias behavior.

²⁴ This author believes that, had the survey evidence been of the more traditional ambiguous kind, the results in the final row of Survey A would have exhibited a stronger Against column response, and a less strong Neutral column response. It is easier to misread conflicting evidence as supportive of your initial hypothesis if it is somewhat ambiguous; if it is not ambiguous, the best you can do while still maintaining your initial beliefs is to disregard the new information.

Table 4 Evidence assimilation and attitude change regression results on entire sample

Survey dependent Var.	A		B	
	Evidence assimilation	Attitude change	Evidence assimilation	Attitude change
Constant	3.524*** (0.275)	-0.457* (0.243)	1.226*** (0.385)	-1.446*** (0.253)
Predisposition	0.246*** (0.078)	0.251*** (0.069)	0.355*** (0.113)	0.272*** (0.074)
R^2	0.067	0.087	0.069	0.091

* Significance at 10% level

*** Significance at 1% level

increase evidence presented in Survey B, but only to a much smaller extent. These respondents also tended to place a lot of their evidence assimilation into the Neutral category, most likely to avoid the cognitive conflict between rational utility maximization and confirmatory bias of the unambiguous survey evidence. In fact, the rational response in Survey B (Against) was chosen only 42% of the time by initial supporters, while the symmetric assimilation category in Survey A (Opposition, Favor) reports a 62% assimilation rate. The bias in Survey B appears to be substantially higher than that reported in Survey A.

Further support for this interpretation of the results can be found from a 2×2 ANOVA, with the dependent variable being evidence assimilation, the first factor being initial predisposition, and the second factor having received Survey A versus Survey B. The results confirm the significance of initial predisposition on evidence assimilation ($F = 5.23$, $df = 5$, $p < 0.001$); they also confirm that Survey A and Survey B results were distinct ($F = 31.7$, $df = 1$, $p < 0.001$), lending weight to the argument that the bias in Survey B is different from that in Survey A. Finally, and as would be expected, the ANOVA results indicate the insignificance of the interaction of the two factors on evidence assimilation ($F = 0.94$, $df = 5$, $p < 0.46$). While initial predisposition and survey type are important in explaining subsequent evidence assimilation, there is no reason to believe that initial predisposition given a particular survey type should have any affect whatsoever.

Attitude polarization, one result of confirmatory bias behavior (Lord et al 1979) and a frequent tool used to confirm its existence, also appears to be in evidence from the results displayed in Table 3.²⁵ The frequencies in Table 3 show that, for both surveys, if a subject's initial predisposition was in support (opposition) of Bush Administration tax policy, then after assimilation of the evidence, the subject updated their attitude towards Bush Administration tax policy in a more (less) favorable light than initially predisposed opponents (supporters) tended to do. A 2×2 ANOVA on attitude change, with initial predisposition ($F = 8.4$, $df = 5$, $p < 0.001$) and survey type ($F = 14.98$, $df = 1$, $p < 0.001$) as the two factors, confirms the significance of these results.

²⁵ These results from Survey A and Survey B are both normally distributed as well.

All of these conclusions, however, would be more persuasive if they were based on results produced from the entire data set. To test for the prevalence of confirmatory bias in the entire subject pool, Table 4 presents the results for Survey A and B of regressions run on the complete set of observations. The independent variable (initial predisposition) is the same for both regressions, and is run on subsequent evidence evaluation and attitude change.

We find that initial predisposition affected subsequent evidence evaluation and attitude change significantly in Survey A, and to an even greater extent (larger coefficients) in Survey B. These persuasive results, based on the entire survey pool, imply that confirmatory bias is widespread enough (or strong enough) to affect aggregate outcomes. If confirmatory bias existed, but only for extreme behavior which was corrected for in an aggregate sense, it would not be of great interest to economists in particular. These results show, however, that the effects of confirmatory behavior can be strong enough to affect aggregate outcomes. In terms of welfare effects, this could imply strong constituent support for tax policies (independent of other policy issues) that may ultimately be utility decreasing.²⁶

A number of criticisms can be raised regarding these results. First, how relevant was this survey to the targeted subject pool? Specifically, are students in this sample actually paying for their schooling (as opposed to receiving grant aid) and so directly affected by this tax proposal evidence? In a similar vein, are these tax policy proposals (of a few thousands dollars) marginally significant to the students, or, are their incomes high enough that these proposals may be perceived as marginally insignificant? Regarding grant aid, it is unlikely that very many (if any) of the subjects participating in this survey were receiving 100% of their full college tuition in grant aid. Records for the 2003 MSU student population show that only 5% of the entire student body had full, non tax deductible grant aid.²⁷ As for annual incomes, the summary statistics presented in Table 1 should dispel any concern over relative irrelevance of the size of the tax policy proposals—92.8% of the survey respondents reported an annual income of \$20,000 or less, and no survey respondents reported an annual income above \$75,000. From anecdotal evidence gathered in the classroom, it also appears that most of these students are paying for some (if not all) of their college education themselves (as opposed to relying on generous parents).

Another criticism of the survey evidence presented in this study may be that order effects are driving the results, rather than any sort of confirmatory bias. Order effects occur when the question order of a survey affects survey responses in a seemingly arbitrary way—if the questions had been presented in a different order, for example, would the respondents have answered differently? Specifically, if the tax policy evidence (Part II) had been presented first, and the initial predisposition

²⁶ Akerlof (1989), in an insightful paper on voting behavior, considers a similar setting when he models decision-making behavior subject to an “illusion” bias. He finds, as we do, that while “it is usually assumed that individuals vote in their own best interest...this best interest may be neither obvious nor *pleasant to contemplate*...the desire and ability for self-delusion can lead to poor social decisions.” (p. 10, italics added).

²⁷ Common Data Set, online at: <http://www.commondataset.org/> (The Office of Planning and Analysis, Montana State University-Bozeman 2003).

questions (Part I) second, would the subjects have answered the questions differently? This assumes, of course, and as it should, that the survey is not determining initial proclivity, but merely measuring it.

Evidence for a limited concern over order effects driving the presented results comes from the seminal, and still well referenced, exposition on survey administration by Dillman (2000), where the author notes that order effects are smaller for self-administered surveys (like this one) than for telephone or instructor administered surveys. Further, Dillman notes that structural manipulations (such as presenting the survey material on different pages and with different headings, as was done here) has been found to reduce order effects substantially. Finally, analysis of the empirical results seems to support a confirmatory bias interpretation of the evidence over an order effects interpretation. If the results had been driven primarily by order effects, then there should not have been any significant difference between the numerical results presented in Survey A and those presented in Survey B. The fact that these two surveys, while symmetric in order construct, produced results that were not symmetric in response, points to a deeper explanation for the evidence. The most convincing explanation, this author believes, is the presence of a confirmatory bias judgment error.^{28,29}

Finally, concern over the empirical relevance of contingent value surveys, as opposed to revealed preference evidence, continues to be a subject of contention amongst many economists. The author readily admits that survey evidence is not as compelling as direct market evidence, but it is difficult to come up with an obvious revealed preference dataset that can accurately test for confirmatory bias behavior; indeed no one has yet done it. Instead, the aim of this paper is to take an initial stab at uncovering economic-oriented examples of the phenomenon since the literature to date lacks such empirical studies.

5 Conclusion

This paper provides an empirical test of the confirmatory bias behavioral phenomenon; manifestation of which implies implacable, even non-optimal decision-making. Two specific hypotheses were put forth: (1) that the confirmatory bias effect can occur when assimilating relatively unambiguous evidence, and (2)

²⁸ A somewhat different interpretation, suggested to the author by Ken Troske, is that conservatives (who tend to support Bush tax policies) are more politically ideological than liberals (who tend not to support Bush tax policy). This difference in political ideological degree may be the driving factor behind the nonsymmetric results, rather than that of prospect theory and varying preferences for gains versus losses. From the survey data collected it is impossible to be certain which interpretation is correct, but either way, both interpretations support the assumption that order effects are not driving the results; they merely offer different explanations for the motivation behind the confirmatory bias effect.

²⁹ An anonymous reviewer also suggested that the nonsymmetric results may be due to a framing effect whereby people “support cuts” more than they “oppose increases” (somewhat similar to the asymmetry of information it takes to confirm versus falsify hypotheses (Trope and Thompson 1997; Cameron and Trope 2004)). While we can not rule out this interpretation entirely either, we believe it unlikely to be the complete explanation for our results, or more than 62% of those initially predisposed against Bush Administration tax policy would have more positively “supported” the tax cuts.

that the confirmatory bias effect is significant enough to affect aggregate outcomes. Evidence was found in support of both of these hypotheses. Further, the survey results indicate that confirmatory bias may be stronger when assimilating utility decreasing evidence (losses) than when assimilating utility increasing evidence (gains). Exploring the motivations and welfare consequences of these results would be a useful direction for future research, as well as testing further for the very existence of confirmatory bias in other applied settings. The initial results presented in this study, if assimilated without a traditional-leaning confirmatory bias, point to the need to tread carefully whenever assuming any rigid interpretation of the rational, neo-classical, utility maximizing economic model.

Acknowledgments Many special thanks to Rob Fleck, Dino Falaschetti, Susan Feigenbaum, Scott Plous, Ron Harstad, David Mandy, Jeff Milyo, Ken Troske, and two anonymous reviewers for their insight, criticism, support, and advice.

Appendix 1: Survey questions

Part I: Initial predisposition

Q1 To what extent do you think the tax cuts which Congress passed and George W. Bush signed into law during this past presidential term have affected the US economy?

- greatly helped the US economy
- somewhat helped the US economy
- had no effect on the US economy
- somewhat hurt the US economy
- greatly hurt the US economy

Q2 In the long run do you think tax cuts strengthen or weaken the economy?

- greatly weaken the economy
- somewhat weaken the economy
- have no effect on the economy
- somewhat strengthen the economy
- greatly strengthen the economy

Q3 In recent years, President Bush and Congress have made major cuts in federal income tax rates. What is your opinion regarding this tax agenda?

- strongly approve
- somewhat approve
- have no opinion
- somewhat disapprove
- strongly disapprove

Q4 Just your opinion: What do you think about the size of President Bush's tax cuts? Should they have been:

- much smaller than they were
- somewhat smaller than they were
- about the same as they were
- somewhat bigger than they were
- much bigger than they were

Part II: Evidence assimilation

Survey A:

President Bush is due to propose to Congress, within the next month, an additional tax cut to be effective for the 2004 tax year. This tax cut takes the form of an extension to the HOPE scholarship Credit. Currently the HOPE Scholarship Credit is an income tax credit for students in the first two years of college (or other eligible post-secondary training), which allows them to credit their tax returns up to 100% of the first \$1,000 of tuition and fees and 50% of the second \$1,000 (the amounts are indexed for inflation after 2001). The credit is available on a per-student basis for net tuition and fees (less grant aid) paid for college enrollment after 31 December 1997.

The proposal President Bush is presenting to Congress next month extends the benefits and terms of this credit. It will be available for up to four years of college (instead of just the first two), and the amount of tuition and fees available for credit will be raised to 100% of the first \$2,000 of tuition and fees and 50% of the second \$2,000 (the amounts are still indexed for inflation). Note that other education and learning credits—such as the Lifetime Learning Credit—will not be affected in any way. If you are a current college student, therefore, you will unambiguously benefit from this proposal.

Survey B:

President Bush is due to propose to Congress, within the next month, a tax increase to be effective for the 2004 tax year. This tax increase takes the form of a reduction of the HOPE scholarship credit. Currently the HOPE Scholarship Credit is an income tax credit for students in the first 2 years of college (or other eligible post-secondary training), which allows them to credit their tax returns up to 100% of the first \$1,000 of tuition and fees and 50% of the second \$1,000 (the amounts are indexed for inflation after 2001). The credit is available on a per-student basis for net tuition and fees (less grant aid) paid for college enrollment after 31 December 1997.

The proposal President Bush is presenting to Congress next month reduces the benefits and terms of this credit. It will only be available for the first year of college (instead of the first two years as before), and the amount of tuition and fees available for credit will be reduced to 50% of the first \$1,000 of tuition and fees and 25% of the second \$1,000 (the amounts are still indexed for inflation). Note that other education and learning credits—such as the Lifetime Learning Credit—will not be affected in any way. If you are a current college student, therefore, you will unambiguously lose from this proposal.

Q5 The HOPE Scholarship Credit is an income tax credit for:

- employee moving expenses
- college tuition and fees expenses
- on the job training expenses
- medical expenses

Q6 The proposed amendment to the HOPE Scholarship Credit, if passed, would take effect:

- retroactively back to tax year 1997
- in tax year 2010
- in tax year 2007
- in tax year 2004

Q7 The proposed amendment to the HOPE Scholarship Credit includes increasing (decreasing) the number of years college students are eligible to claim this credit from two to:

- one
- six
- four
- three

Q8 Please indicate your support or opposition for this latest tax proposal:

- strongly support
- somewhat support
- have no opinion
- somewhat oppose
- strongly oppose

Part III: Attitude polarization

Q9 After reading and thinking about the new proposed tax amendment, which of the following statements comes closest to your point of view?

- I now have much greater faith in President Bush's tax agenda
- I now have greater faith in President Bush's tax agenda

- I now have less faith in President Bush's tax agenda
- I now have much less faith in President Bush's tax agenda
- My opinion has not changed

Q10 After reading and thinking about the new proposed tax amendment, which of the following statements comes closest to your point of view?

- I am much less inclined to support President Bush's tax agenda
- I am somewhat less inclined to support President Bush's tax agenda
- I am somewhat more inclined to support President Bush's tax agenda
- I am much more inclined to support President Bush's tax agenda
- My opinion has not changed

Q11 After reading and thinking about the new proposed tax amendment, which of the following statements comes closest to your point of view?

- I see passage of all of President Bush's income tax proposals much more likely
- I see passage of all of President Bush's income tax proposals somewhat more likely
- I see passage of all of President Bush's income tax proposals somewhat less likely
- I see passage of all of President Bush's income tax proposals much less likely
- My opinion has not changed

References

- Akerlof GA (1989) The economics of illusion. *Econ Pol* 1(1):1–15
- Arkes HR, Wartmann RL, Saville PD, Harkness AR (1981) Hindsight bias among physicians weighing the likelihood of diagnoses. *J Appl Psychol* 66(2):252–254
- Ausubel LM (1991) The failure of competition in the credit card market. *Am Econ Rev* 81(1):50–81
- Babcock L, Loewenstein G (1997) Explaining bargaining impasse: the role of self-serving biases. *J Econ Perspect* 11(1):109–126
- Barber BM, Odean T (2001) Boys will be boys: gender, overconfidence, and common stock investment. *Q J Econ* 116(1):261–292
- Baumann AO, Deber RB, Thompson GG (1991) Overconfidence among physicians and nurses: the 'Micro-Certainty, Macro-Uncertainty' phenomenon. *Soc Sci Med* 32(2):167–174
- Borum R, Otto R, Golding S (1993) Improving clinical judgment and decision making in forensic evaluation. *J Psychol Law* 21(1):35–76
- Bruner JS, Potter MC (1964) Interference in visual recognition. *Science* 144(3617):424–425
- Camerer CF (1997) Progress in behavioral game theory. *J Econ Perspect* 11(4):167–188
- Camerer CF, Lovallo D (1999) Overconfidence and excess entry: an experimental approach. *Am Econ Rev* 89(1):306–318
- Cameron JA, Trope Y (2004) Stereotype-biased search and processing of information about group members. *Soc Cogn* 22(6):650–672
- Cross P (1977) Not can but will college teaching be improved. *New Dir High Ed* 17:1–15
- Dardenne B, Leyens JP (1995) Confirmation bias as a social skill. *Pers Soc Psychol Bull* 21(11):1229–1239

- Darley JM, Gross PH (1983) A hypothesis-confirming bias in labeling effects. *J Pers Soc Psychol* 44(1):20–33
- Dillman DA (2000) *Mail and internet surveys: the tailored design method*, 2nd edn, Wiley, New York
- Dougherty TW, Turban DB, Callender JC (1994) Confirming first impressions in the employment interview: a field study of interviewer behavior. *J Appl Psychol* 79(5):659–665
- Downs A (1957) *An economic theory of democracy*. Harper and Row Publishers, New York
- Dubra J (2004) Optimism and overconfidence in search. *Rev Econ Dyn* 7:198–218
- Dunning D, Meyerowitz JA, Holzberg AD (1989) Ambiguity and self-evaluation: the role of idiosyncratic trait definitions in self-serving assessments of ability. *J Pers Soc Psychol* 57(6):1082–1090
- Eales JS, Engel BK, Hauser RJ, Thompson SR (1990) Grain price expectations of Illinois farmers and grain merchandisers. *Am J Agric Econ* 72(3):701–708
- Ganguly AR, Kagel JH, Moser DV (2000) Do asset market prices reflect traders' judgment biases? *J Risk Uncertain* 20(3):219–245
- Golec J, Tamarkin M (1995) Do betters prefer long shots because they are risk-lovers, or are they just overconfident? *J Risk Uncertain* 11:51–64
- Griffin D, Tversky A (1992) The weighing of evidence and the determinants of confidence. *Cognit Psychol* 24(3):411–435
- Haverkamp BE (1993) Confirmatory bias in hypothesis testing for client-identified and counselor self-generated hypotheses. *J Couns Psychol* 40(3):303–315
- Hvide HK (2002) Pragmatic beliefs and overconfidence. *J Econ Behav Organ* 48:15–28
- Kahneman D, Tversky A (1973) On the psychology of prediction. *Psychol Rev* 80(4):237–251
- Kahneman D, Slovic P, Tversky A (eds) (1982) *Judgment under uncertainty: heuristics and biases*. Cambridge University Press, New York
- Keren G (1987) Facing uncertainty in the game of bridge: a calibration study. *Organ Behav Hum Decis Process* 39(1):98–114
- Keren G (1988) On the ability of monitoring non-veridical perceptions and uncertain knowledge: some calibration studies. *Acta Psychol* 67(2):95–119
- Klayman J, Ha YW (1987) Confirmation, disconfirmation, and information in hypothesis testing. *Psychol Rev* 94(2):211–228
- Libecap GD, Hansen ZK (2000) 'Rain Follows the Plow' and dryfarming doctrine: the climate information problem and homestead failure in the Upper Great Plains, 1890–1925. National Bureau of Economic Research Working Paper, Historical Paper 127
- Lord CG, Ross L, Lepper MR (1979) Biased assimilation and attitude polarization: the effect of prior theories on subsequently considered evidence. *J Pers Soc Psychol* 37(11):2098–2109
- Lord CG, Lepper MR, Preston E (1984) Considering the opposite: a corrective strategy for social judgment. *J Pers Soc Psychol* 47(6):1231–1243
- Madsen JB (1994) Tests of rationality versus an 'Over Optimist' bias. *J Econ Psychol* 15:587–599
- Miller AG, McHoskey JW, Bane CM, Dowd TG (1993) The attitude polarization phenomenon: role of response measure, attitude extremity, and behavioral consequences of reported attitude change. *J Pers Soc Psychol* 64(4):561–574
- Nelson P (1994) Voting and imitative behavior. *Econ Inq* 32:92–102
- Nickerson RS (1998) Confirmation bias: a ubiquitous phenomenon in many guises. *Rev Gen Psychol* 2(2):175–220
- Nisbett RE, Ross L (1980) *Human inference: strategies and shortcomings of social judgment*. Prentice Hall, New Jersey
- The Office of Planning and Analysis, Montana State University-Bozeman (2003) *Common Data Set*. The Thomson Corporation and US News and World Report
- Peltzman S (1980) The growth of government. *J Law Econ* 23(2):209–287
- Plous S (1991) Biases in the assimilation of technological breakdowns: do accidents make us safer? *J Appl Soc Psychol* 21(13):1058–1082
- Preston CE, Harris S (1965) Psychology of drivers in traffic accidents. *J Appl Psychol* 49(4):284–288
- Rabin M (1998) Psychology and economics. *J Econ Lit* 36(1):11–46
- Rabin M, Schrag J (1999) First impressions matter: a model of confirmatory bias. *Q J Econ* 114(1):37–82
- Snyder M, Swann WB (1978) Behavioral confirmation in social interaction: from social perception to social reality. *J Exp Soc Psychol* 14:148–162
- Strohmer DC, Shivy VA (1994) Bias in counselor hypothesis testing: testing the robustness of counselor confirmatory bias. *J Couns Dev* 73:191–197

-
- Svenson O (1981) Are we all less risky and more skillful than our fellow driver? *Acta Psychol* 47:233–250
- Swann WB, Snyder M (1980) On translating beliefs into action: theories of ability and their application in an institutional setting. *J Pers Soc Psychol* 38(6):879–888
- Taylor SE, Brown JD (1988) Illusion and well-being: a social psychological perspective on mental health. *Psychol Bull* 103(2):193–210
- Trope Y, Mackie DM (1987) Sensitivity to alternatives in social hypothesis-testing. *J Exp Soc Psychol* 23:445–459
- Trope Y, Thompson EP (1997) Looking for truth in all the wrong places? Asymmetric search of individuating information about stereotyped group members. *J Pers Soc Psychol* 73(2):229–241
- Tweney RD, Doherty ME, Worner WJ, Pliske DB, Mynatt CR (1980) Strategies of rule discovery in an inference task. *Q J Exp Psychol* 32:109–123
- Wason PC (1960) On the failure to eliminate hypotheses in a conceptual task. *Q J Exp Psychol* 12:129–140
- Watson A, Rodgers T, Dudek D (1998) The human nature of management consulting: judgment and expertise. *Manage Dec Econ* 19(7–8):495–503