Psychological Constraints on Transparency in Legal and Government Decision Making[1]

Robert J. MacCoun, Richard and Rhoda Goldman School of Public Policy, University of California, Berkeley; E-Mail: maccoun@berkeley.edu

In a democratic society, the desirability of openness and transparency in government decision making would seem nearly self-evident. The average newspaper reader has consumed a steady diet of examples of deception, bigotry, cronyism, and corruption by public officials, providing obvious arguments in favor of greater transparency. But my goal for this essay is to illustrate ways in which the complexity of cognition and motivation make transparency difficult to achieve, even with the full cooperation of well-intentioned government actors. The claim is not that particular judgments themselves are complex; indeed many seeming

[1] An early version was presented at the International Workshop on Publicity/Transparency in Democratic Decision Making, Univ. cath. de Louvain, Chaire Hoover d'ethique economique et sociale, Louvain-la-Neuve, Belgium, 23 May 2005. I thank Gene Bardach, Axel Gosseries, David Kirp, Alain Marciano, Daniel Naurin, Patricia Poelier, Alasdair Roberts, David Stasavage and an anonymous referee for helpful comments.

complex judgments are actually well explained by simple linear models. Rather, the claim is that the complexity of the underlying machinery of the brain makes it difficult for actors to either consciously monitor or control their judgment process (Wegner and Bargh 1998). Moreover, some attempts to achieve transparency may have unintended and undesirable consequences.

Any judgments about the merits of transparency, and optimal design of transparency mechanisms must necessarily involve a host of complex issues involving moral and political philosophy, political science, economics, history, and sociology. Many of these perspectives are addressed elsewhere in this symposium. Also, the psychological literature I draw upon relies heavily on experimentation with students. These studies provide strong causal inference about cognitive mechanisms, but weak external validity. Importantly, these experiments inevitably exclude organizational structure, historical context, and most of the tactical and dramatic elements that characterize the "rough and tumble" of real-world politics. Thus, the ideas presented here should be considered hypotheses for further consideration rather than firm conclusions about actual political systems. Fortunately, a number of scholars have begun to investigate some of the social psychological aspects of transparency and political deliberation in legislative bodies, citizen focus groups, political internet chat rooms, and other settings (e.g. Bächtiger, Spörndli, Steenbergen and Steiner 2005; Conover and Searing 2005; Steenbergen, Bächtiger, Spörndli and Steiner 2004; also see Rosenberg 2005).

The Limits of Introspection

There is broad agreement among contemporary psychologists that people are generally unable to fully perceive, and reliably and validly report, the causes of their behavior. The Freudian account of repression or other defense mechanisms is familiar to most readers. But even in the absence of dark subterranean motives, our span of awareness permits only limited access to most other cognitive processes (Hassin, Uleman and Bargh 2005; Wegner and Bargh 1998; Wilson 2002). In a classic 1977 article, Nisbett and Wilson argued that when people attempt to explain their behavior, "they do not do so on the basis of any true introspection". They showed that research participants repeatedly failed to detect experimental factors that were demonstrably influencing their behavior. At the same time, research participants routinely cited "causes" for their behavior that were in fact demonstrably uncorrelated with their responses. Nisbett and Wilson argued that their participants' "explanations" were based not on introspective access, but rather on *a priori*, implicit causal theories about whether a particular stimulus seems plausible as an account of their behavior. In other words, self-attributions reflect not direct perception of the causes of one's behavior, but rather "common sense" or "lay theories" about what those causes might be.

Evidence for limited introspective access runs counter to "folk psychological" notions of mind and rationality, but is entirely in accordance with contemporary scientific psychological theory and research – much of it developed in the years after Nisbett and Wilson's analysis. There is now a considerable body of converging evidence (from research on learning, memory, stereotyping, attitude change, and perception) the brain has two distinct modes of information processing, often working in tandem. Table 1 summarizes the basic properties of System 1 (or associative) and System 2 (or rule-based) cognition (Chaiken and Trope 1999; Sloman 1996; Stanovich 2004). Of particular relevance to the question of transparency, System 2 is conscious and deliberative, but System 1 is characterized by automatized and/or implicit cognition. Automatic cognitive processes can be triggered without conscious effort or intent, are difficult to interrupt without disruption, and often run to completion without any apparent awareness by the actor (Wegner and Bargh 1998)

The distinction between System 1 and 2 modes of cognition suggests plausible constraints on the transparency of decision processes. In particular, *System 1 processing is inherently lacking in transparency – for the decision maker as well as the audience.* Can we simply require decision makers to engage in System 2 processing? Unfortunately, decision makers can simply choose which cognitive processes to engage in. System 2 processes rely heavily on a substrate of System 1 products, though some tasks are more deliberative than others. For example, the process of grading an arithmetic test is inherently more transparent than the process of grading an essay exam; the process of choosing the best credit card is inherently more transparent than the process of choosing a painting to display on your wall. Few important legal or policy decisions involve straightforwardly algorithmic reasoning. Arguably, System 2 thinking might be a curvilinear function of familiarity; very novel tasks demand more System 2 thinking because they can't (yet) be routinized, and very familiar repetitive tasks may be more likely to have produced formal rules. It is the domain in the middle – tasks that are somewhat familiar but not formulaic – that are likely to be most opaque in their processing.

Table 1: *Properties of the two systems (adapted from Stanovich 2004)*

System 1	*System 2*
Associative	Rule-based
Holistic	Analytic
Parallel	Serial
Automatic	Controlled
Relatively effortless	Relatively effortful
Relatively fast	Relatively slow

Due to the limits of introspection, cognitive processes are sometimes more complex than the actor realizes. But it is sometimes the case that actors cite a complex configuration of factors underlying their judgments, when in fact their judgments are well described by a simple linear model, in which a few variables are weighted and then added together (Dawes, Faust and Meehl 1989). A classic example is Ebbesen and Konecni's (1975) study of bail setting decisions by judges. In reactions to fictitious cases, judges made use of a number of legally relevant criteria, including the accused's ties to the area and prior criminal record. But a regression analysis of actual bail decisions failed to reveal any influence of these factors; instead, judges relied almost exclusively on the recommendations made by the district attorney – in essence, acting as a "rubber stamp".

Unintended Consequences of Promoting Transparency

Having argued for the psychological difficulty of achieving transparency, I now list reasons why attempts to encourage transparency can go awry. Psychologists have a seemingly endless supply of biases and perversities to invoke against optimal rationality, but I limit the focus to a few well documented phenomena arising in circumstances that might be expected to promote transparency. These phenomena by no means imply that the net or total effect of an intervention is undesirable, which would require a broader assessment of costs and benefits in context. Instead, the argument is just that the interventions are likely to perform less well than intended.

Accountability Can Have Perverse Effects

Most efforts at transparency lead decision makers to expect some degree of accountability to an audience. In the face of accountability pressures, Tetlock (2002) argues that people are more like "intuitive politicians" than like intuitive scientists. He and his students have shown that the effects of expecting to be held accountable for a decision vary depending on whether the audience's views are known or unknown, and the nature of those views (see Lerner and Tetlock 1999). If decision makers expect to be accountability to an audience of unknown views, they engage in "pre-emptive self-criticism" – a process of careful consideration of alternatives and systematic reasoning about evidence. But if the audience views are known, the decision maker is more likely to engage in some "attitude shifting" – moving in the direction of the audience's viewpoint.[2]

[2] Tetlock also considers what happens when one only discovers that one is accountable after a decision has already been rendered. In such situations, accountability cannot affect the decision but, if the audience is dissatisfied, the decision maker may engage in "defensive bolstering" in an effort to rationalize the position.

An experimental study of the effects of juror anonymity can be interpreted in this light. Hazelwood and Brigham (1998) compared verdicts rendered by 20 anonymous and 20 non-anonymous juries in a student disciplinary case. When the case against the student defendant was strong, the anonymous juries were considerably more likely to vote for conviction and punishment, suggesting that the non-anonymous juries were reluctant to publicly sanction a fellow student.

Another means of promoting transparency is to ask decision makers to self disclose any personal biases or conflicts of interest. But there are a number of problems with this idea. First, it assumes people are aware of their biases, but in fact most people believe they are less biased than others (Pronin, Gilovich and Ross 2004). Second, Cain, Loewenstein and Moore (2005) have demonstrated that disclosures can actually increase the influence of bias in expert advice; apparently experts – having "come clean" feel they have full license to say what they want without caution or qualification. Making matters worse, their audience – having heard the expert come clean – no longer adjust their understanding of the advice to correct for possible bias from the source.

Group Processes Can Amplify Bias and Discourage Deliberation

Group deliberation is commonly endorsed as a means of promoting legitimacy and encourage a diversity of viewpoints to be heard. Of course, aggregation will cancel out random error (noise). And it can cancel out systematic error (bias) provided that some biases offset others. But what about shared biases? In general, group deliberation tends to *attenuate* shared individual biases when a "wrong" (or biased) response is readily recognizable as such by the group. But in judgmental tasks lacking a shared conceptual scheme for recognizing a correct response, deliberation tends to *amplify* bias (Kerr, MacCoun and Kramer 1996).

In collective decision making, transparency is often sought by requiring decision makers to state their views publicly (rather than privately). A basic finding in social psychology is that public commitment to a position makes people more resistant to moderating their views in light of subsequent argument (Jellison and Mills 1969). In criminal juries, for example, premature voicing of opinions can make deliberation "verdict-driven" rather than "evidence-driven" (Hastie, Penrod and Pennington 1983). Norbert Kerr and I found that the instruction to conduct open "show-of-hands" polling led to a significantly larger number of "hung" (deadlocked) juries than did secret balloting (Kerr and MacCoun 1985). Gosseries (2005) quotes James Madison's concern about this kind of premature commitment at the 1787 American Constitutional Convention: "Had the members committed themselves publicly at first, they would have afterwards supposed consistency required them to maintain their ground, whereas by secret discussion no man felt himself obliged to retain his opinions any longer than he was satisfied of their propriety and truth, and was open to the force of argument".

Gosseries (2005) argues that deliberating actors need to be able to "try out ideas out of the blue with the risk of having to abandon them straightaway (trial and error), to show hesitation, to reconsider the issues again and again with a fresh eye, the actual deliberation may not be more than the juxtaposition of pre-prepared statements with no actual interaction taking place". In the organizational behavior literature, this concern is promoted with the concept of "brainstorming", in which group members are encouraged to generate a large body of candidate solutions to a problem while postponing any criticism or reality testing. Unfortunately, in practice, brainstorming groups tend to prematurely seize upon early plausible candidates; a greater pool of candidate solutions can be generated by distributing the ideas generated by individuals working alone (Mullen, Johnson and Salas 1991).

Being Explicit Can Distort Goals and the Willingness to Make Tradeoffs

The System 1 processes described earlier presumably evolved largely to permit fast responding and to achieve metabolic efficiency – extended bouts of highly controlled conscious processing are very taxing, as exam takers and novice meditators can attest. But another advantage of System 1 processes is that parallel distributed processing facilitates the simultaneous pursuit of multiple goals or objectives. I'm not referring here to multitasking (doing many different tasks at once), but trying to achieve multiple ends with the same task. Robbennolt, Darley and I (2003) have argued that legal decision making routinely requires fact finders to pursue multiple goals; e.g. consistency with the evidence, allocation proportional to fault, compensation proportional to need, retribution for or deterrence of negligent or egregious conduct and so on. We described four basic principles of cognitive goal management:

> First, the *principle of equifinality* holds that some goals may be alternately satisfied through multiple pathways. …Second, *the principle of best fit* holds that pathways may sometimes better fulfill some goals than others. …Third, *the principle of multifinality* holds that a particular pathway may accomplish multiple goals simultaneously. …Finally, the principle of *goal incompatibility* holds that some objectives will inevitably conflict and, thus, be difficult or impossible to satisfy concurrently.

We concluded that parallel constraint satisfaction networks – a type of System 1 cognitive architecture – are particularly well suited to achieving these four principles.

Interestingly, our proposal resonates well with Sternberg's (1998) "balance theory of wisdom". Sternberg's theory defines wisdom as "the application of tacit knowledge as mediated by values toward the achievement of a common

good through a balance among multiple (a) intrapersonal, (b) interpersonal and (c) extrapersonal interests in order to achieve a balance among (a) adaptation to existing environments, (b) shaping of existing environments and (c) selection of new environments".

A concern with efforts to mandate transparency is that the explicit listing of reasons tends to valorize certain easily articulated or salient goals at the expense of others. This is a good thing if the goals that get neglected are socially inappropriate (e.g. pecuniary conflicts of interest), but it is a bad thing if it discourages the practice of "wisdom" in Sternberg's sense. In a clever set of studies, Timothy Wilson and his collaborators have demonstrated that process of listing one's reasons for a judgment can actually produce judgments that are poorer by a variety of criteria (see Wilson 2002). For example, they have shown that analyzing one's reasons:

- makes lay participants' judgments become less correlated with those of experts,
- makes people less satisfied with their choice at a later time,
- makes people more likely to later change their mind about the judgment they reached, and
- produces greater overconfidence, greater reliance on confirmatory bias and reduces the correlation between one's announced judgment and one's subsequent behavior.

Wilson argues that analyzing reasons leads people to give disproportionate weight to those factors that are salient or easily brought to mind, relative to other factors that otherwise would have (and often should have) weighed heavily in their judgment process. A possible exception to this general pattern is that an explicit "reasons analysis" may be more helpful in situations where the decision maker has available an explicit algorithm or conceptual scheme which, if used, will improve the chance of finding the right answer (McMackin and Slovic 2000).

In schools of public policy analysis, we actively instill our students with the belief that good policy making requires a willingness to make difficult tradeoffs.[3] Unfortunately, people are often highly resistant to confronting such tradeoffs. A great many of my students – even Masters in Public Policy candidates with several years of government experience under their belts – display visceral discomfort when confronted with expert analyses of the economic valuation of human life. Many students roll their eyes, shake their heads, or scowl. When asked to explain, they struggle to verbalize their feelings; they find it distasteful to place a value on human life, but they can't say why, and most acknowledge the need for policy analysts to do so.

[3] The discussion that follows is adapted from MacCoun (2000).

A psychological analysis of this discomfort is provided by the "taboo tradeoffs" theory of Fiske and Tetlock (1997). Drawing on Fiske's earlier work, the authors contend that relations in all societies are governed by various combinations of four fundamental psychological templates: We sometimes categorize individuals and treat category members identically (communal sharing), we sometimes treat individuals by their rank within a group (authority ranking), we sometimes keep score of outcomes and strive to equalize them (equality matching), and we sometimes value outcomes on an absolute metric and make tradeoffs among them (market pricing). Each template has its own rules of appro-priate conduct, its own norms of distributive fairness, and most crucially, its own consensually agreed upon domains of operation in a community's life.

Fiske and Tetlock (1997: 294) argue that "cost-benefit analysis ignores and usually does violence to normative distinctions that people value as ends in themselves". They recognize the normative value of formal cost-benefit analysis; they recognize that "taboo tradeoffs are unavoidable. …In practice, there is a limit to the dollars we will spend to enhance our own personal safety at the workplace or in cars or airplanes, and we will certainly spend less for the safety of others" (Fiske and Tetlock 1997: 290). But they argue that attempts to apply market pricing to the domain of human life will inevitably encounter resistance: "It is gauche, embarrassing, or offensive to make explicit trade-offs among the concurrently operative relational modes" (Fiske and Tetlock 1997: 273). This suggests that explicitness can discourage tradeoffs that we might otherwise prefer to make. Though tradeoffs can be inappropriate or corrupt, most policy analysts believe some such tradeoffs are a necessary part of responsible governance.

Another example of this kind of tension is Tribe's (1971) argument that standards of proof at trial should not be explicitly quantified. He contends that an explicitly probabilistic criterion (e.g. "you must be more than 95 percent certain") is undesirable because it provides an explicit statement of error rates; thus, certainty less than 100 percent implies that some innocent defendants might be acquitted. Tribe worries that this admission might delegitimize the system. But Kagehiro (1990) has shown that jurors actually perform better with quantified standards of proof than verbal standards; thus, the unwillingness of courts to openly acknowledge error tradeoffs may increase legal error.

There are also many controversial examples involving vice policy. For example, "harm reduction" interventions like needle exchange are demonstrably beneficial from a public health standpoint, but when they are too visibly discussed and acknowledged, they raise objections that the government is "sending the wrong message", tacitly endorsing injection drug use (MacCoun and Reuter 2001). The Dutch are notorious for comfortably tolerating policies that are difficult to reconcile when examined explicitly. For example, the Netherlands prohibits cannabis possession, yet has a written policy of non-enforcement of this

law for small quantities; this allows them to target high level traffickers, to avoid punishing low level offenders, to separate the cannabis and hard drug markets, to control retail vendor behavior, and reduce the "forbidden fruit" effect of banning marijuana (MacCoun and Reuter 2001).

A related reason why public accounts may differ from private reasoning is that ex ante, declared rules are intended to have an expressive, injunctive function. But ex post, holding citizens to those standards may sometimes seem unreasonable. Legal scholars have drawn a distinction between conduct rules (which tell the public how to behave) and decision rules (which tell judges how to evaluate the public's conduct). Dan-Cohen (1984) has examined the benefits and difficulties of achieving an "acoustic separation" between these two types of rules.

Are There Better Ways to Achieve Transparency?

In this essay, I've "problematized" transparency without offering better solutions, but I can make some tentative suggestions about potentially more effective ways of achieving the same ends.

A number of procedural tools for encouraging unbiased decision making have been identified (see MacCoun 1998, 2005), including task decomposition (e.g. special interrogatories for juries), the use of devil's advocacy, and bifurcated decision processes. And there are a variety of less intrusive methods that social scientists routinely use to study and evaluate decision processes, including observational coding, content analysis, process tracing, econometric analysis, quasi-experimental program evaluation designs. Social auditing studies can identify biases using experimentally doctored "cases" that are submitted to decision makers without their knowledge. For example, Ayres (2001) and others have documented racial discrimination in automobile sales, bail bonding, bank lending, and other markets by having white or African American actors with the same personal backgrounds (education, income, etc.) pretend to seek services.

More generally, I offer three conjectures about the likelihood of success in achieving transparency:

1. It may be hardest to achieve transparency for decisions that are intermediate between the very novel (which are more likely to be made in a deliberate fashion) and the very routine (where formal procedures or rules may have been established).
2. At least with respect to psychological constraints, it should be harder to achieve transparency in the reasoning process than for the inputs to and outputs from that process. In order to get inside the black box of the *reasoning process*, we should not expect much help from the decision makers themselves. Indeed, efforts to direct influence their openness can have unintended consequences.

3. Decision making performance ranges from the corrupt and inept to the wise and inspired. Efforts to increase transparency can and probably do eliminate many decisions from the worst end of the continuum, but it is conceivable that they do so at the expense of impairing high quality decisions at the other extreme.

References

Ayres, Ian (2001). *Pervasive Prejudice? Non-Traditional Evidence of Race and Gender Discrimination*. Chicago: University of Chicago Press.

Bächtiger, André, Markus Spörndli, Marco R. Steenbergen and Jürg Steiner (2005). "The deliberative dimensions of legislatures", *Acta Politica* 40(2): 225-238.

Cain, Daylian M., George Loewenstein and Don A. Moore (2005). "Coming clean but playing dirtier: The shortcomings of disclosure as a solution to conflicts of interest", in Don A. Moore, Daylian M. Cain, George Loewenstein, and Max Bazerman (eds.). *Conflicts of interest: Problems and solutions from law, medicine and organizational settings*. London: Cambridge University Press.

Chaiken, Shelly, and Yaacov Trope (eds.) (1999). *Dual-Process Theories in Social Psychology*. New York: Guilford Press.

Conover, Pamela Johnston and Donald D. Searing (2005). "Studying 'everyday political talk' in the deliberative system", *Acta Politica* 40(3): 269-283.

Dan-Cohen, Meir (1984). "Decision Rules and Conduct rules: On Acoustic Separation in Criminal Law", *Harvard Law Review* 97(3): 625-677.

Dawes, Robyn M., David Faust and Paul E. Meehl (1989). "Clinical versus actuarial judgment", *Science* 243(4899): 1668-1674.

Ebbesen, Ebbe B. and Vladimir J. Konecni (1975). "Decision making and information integration in the courts: The setting of bail", *Journal of Personality and Social Psychology* 32(5): 805-821.

Fiske, Alan Page and Philip E. Tetlock (1997). "Taboo tradeoffs: Reactions to transactions that transgress the spheres of justice", *Political Psychology* 18(2): 255-297.

Gosseries, Axel (2005). "Publicity". *Stanford Encyclopedia of Philosophy*. URL: http://plato.stanford.edu/entries/publicity/

Hassin, Ran R., James S. Uleman and John A. Bargh (eds.) (2005). *The New Unconscious*. Oxford: Oxford University Press.

Hastie, Reid, Steven Penrod and Nancy Pennington (1983). *Inside the Jury*. Cambridge, Massachusetts: Harvard University Press.

Hazelwood, D. Lynn and John C.Brigham (1998). "The effects of juror anonymity on jury verdicts", *Law and Human Behavior* 22(6): 695-713.

Jellison, Jerald M. and Judson Mills (1969). "Effect of public commitment upon opinions", *Journal of Experimental Social Psychology* 5(3): 340-346.

Kagehiro, Dorothy K. (1990). "Defining the standard of proof in jury instructions", *Psychological Science* 1(3): 194-200.

Kerr, Norbert L. and Robert J. MacCoun (1985). "The effects of jury size and polling method on the process and product of jury deliberation", *Journal of Personality and Social Psychology* 48(2): 349-363.

Kerr, Norbert, Robert J. MacCoun and Geoff Kramer (1996). "Bias in judgment: Comparing individuals and groups", *Psychological Review* 103(4): 687-719.

Lerner, Jennifer S. and Philip E. Tetlock (1999). "Accounting for the effects of accountability", *Psychological Bulletin* 125(2): 255-275.

MacCoun, Robert J. (1998). "Biases in the interpretation and use of research results", *Annual Review of Psychology* 49: 259-287.

MacCoun, Robert J. (2000). "The costs and benefits of letting juries punish corporations: Comment on Viscusi", *Stanford Law Review* 52(6): 1821-1828.

MacCoun, Robert. (2005). "Conflicts of interest in public policy research." In Moore, D. A., Cain, D. M., Loewenstein, G. and Bazerman, M. (eds.). *Conflicts of interest: Problems and solutions from law, medicine and organizational settings.* New York: Cambridge University Press.

MacCoun, Robert and Peter Reuter (2001). *Drug War Heresies: Learning from Other Vices, Times, And Places.* New York: Cambridge University Press.

McMackin, John and Paul Slovic (2000). "When does explicit justification impair decision making?", *Applied Cognitive Psychology* 14(6): 527-541.

Mullen, Brian, Craig Johnson and Eduardo Salas (1991). "Productivity loss in brainstorming groups: A meta-analytic integration", *Basic and Applied Social Psychology* 12(1): 3-23.

Nisbett, Richard C. and Timothy D. Wilson (1977). "Telling more than we can know: Verbal reports on mental processes", *Psychological Review* 84(3): 231.

Pronin, Emily, Thomas Gilovich and Lee Ross (2004). "Objectivity in the eye of the beholder: Divergent perceptions of bias in self versus others", *Psychological Review* 111(3): 781-799.

Robbennolt, J. K., J. M. Darley and R. J. MacCoun (2003). "Symbolism and incommensurability in civil sanctioning: Decision-makers as goal managers", *Brooklyn Law Review* 68(3): 1121-1158.

Rosenberg, Shawn (2005). "The empirical study of deliberative democracy: Setting a research agenda", *Acta Politica* 40(2): 212-224.

Sloman, Steven A. (1996). "The empirical case for two systems of reasoning", *Psychological Bulletin* 119(1): 3-22.

Stanovich, Keith E. (2004). *The Robot's Rebellion: Finding Meaning in the Age of Darwin.* Chicago: University of Chicago Press.

Steenbergen, Marco R., André Bächtiger, Markus Spörndli and Jürg Steiner (2004). "Toward a political psychology of deliberation", *Paper presented at the Conference on "Empirical Approaches to Deliberative Politics.* Firenze: European University Institute, .

Sternberg, Robert J. (1998). "A balance theory of wisdom", *Review of General Psychology* 2(4): 347-365.

Tetlock, Philip E. (2002). "Social functionalist frameworks for judgment and choice: Intuitive politicians, theologians, and prosecutors", *Psychological Review* 109(3): 451-471.

Tribe, Lawrence H. (1971). "Trial by mathematics: Precision and ritual in the legal process", *Harvard Law Review* 84(2): 1329-1351.

Wegner, Daniel M. and John A. Bargh (1998). "Control and automaticity in social life", in Daniel Gilbert, Susan T. Fiske and Gardner Lindzey (eds.). *Handbook of social psychology*. 4th edition, Vol. 1 New York: McGraw-Hill, pp. 446-496.

Wilson, Timothy D. (2002). *Strangers to Ourselves: Discovering the Adaptive Unconscious.* Cambridge: Harvard University Press.