

Formal and Effective Autonomy in Health Care

Abraham Schwab
6258 N Wayne #1
Chicago, IL 60660
abe.schwab@ama-assn.org

Senior Fellow
American Medical Association
Ethics Standards Group

KEYWORDS: Autonomy, decision-making, biases, psychology, bounded cognition

ABSTRACT

This essay lays the groundwork for a novel conception of autonomy that I call “effective autonomy”—a conception designed to be genuinely action-guiding in bioethics. As empirical psychology research on the heuristics and biases approach illustrates, decision-making commonly fails to correspond to our desires because of the biases arising from bounded cognition. An individual who is classified as autonomous on contemporary philosophical accounts may fail to be *effectively* autonomous because their decisions are uncoupled from their autonomous desires. Accordingly, continuing attempts to value patient autonomy must go beyond existing philosophical conceptions of autonomy to address the background conditions of human decision-making.

FORMAL AND EFFECTIVE AUTONOMY IN HEALTH CARE

Obstacles to good decision-making by patients are a serious concern in health care. Certain obstacles are generally expected to arise: a lack of information, mental disability, dementia, or immaturity. However, obstacles also arise from unexpected sources. The following example comes from a study of hypothetical preferences regarding medical treatments.[1] In this study, respondents were presented with information about the outcomes of two treatments for lung cancer. Although the statistics presented were identical, they were framed in opposing ways—in terms of survival rates for some respondents and in terms of mortality rates for others (the following is presented exactly as it was presented to participants in the study).

Survival framing:

Surgery: Of 100 people having surgery 90 live through the post-operative period, 68 are alive at the end of the first year and 34 are alive at the end of five years.

Radiation Therapy: Of 100 people having radiation therapy all live through the treatment, 77 are alive at the end of one year and 22 are alive at the end of five years.

Mortality framing:

Surgery: Of 100 people having surgery 10 die during surgery or the post-operative period, 32 die by the end of the first year and 66 die by the end of five years.

Radiation Therapy: Of 100 people having radiation therapy, none die during treatment, 23 die by the end of one year and 78 die by the end of five years.

According to the principle of invariance in rational choice theory, the rationality of a choice is not affected by the framing of the information—the same statistics presented in different ways should result in the same decision. In this case, however, the number of respondents who favored radiation therapy went from 18% for those presented with the survival framing to 44% for those presented with the mortality framing.

In this paper, I will identify three implications of this study and other research like it for the practice of autonomy in health care. First, I will show how the biases of human decision-making have a significant impact on autonomous decision-making. Though the desires or interests of a particular individual may be autonomous, the heuristics and biases of human decision-making may undermine attempts to achieve those desires or interests. Second, this leads me to illustrate how existing procedural and substantive accounts of autonomy (which I call formal) do not address these biases, and I suggest a replacement or modified account that I call “effective autonomy.” My account of effective autonomy, which explicitly addresses the effect of cognitive biases in human decision-making, most resembles the formal accounts of autonomy dubbed as “procedural,” though I will briefly question the value and validity of distinguishing between procedural and substantive accounts. Third and finally, I will offer speculative

policy changes for health care that follow from “effective autonomy.” These will be necessarily brief—more research is required to produce anything more than speculation about needed changes.

As I indicate above, “effective autonomy” is the matching of formally autonomous interests or desires with decisions that will achieve those interests or desires. The term “effective autonomy” is already prevalent in public discourse, though not in philosophy or bioethics. It’s most common use is in political discussions, and is often used to refer to practical autonomy that is not legally recognized. For example, when a group is self-governed despite legal governance by some other body, (i.e. the Kurds in Northern Iraq) they are said to have “effective autonomy.” In my use of effective autonomy in health care, I am suggesting the opposite situation—we already recognize the legal right to autonomy in the medical encounter, but need to address the biases of decision making that may undermine effective autonomy.

Formal Accounts of Autonomy

Most regard autonomy as something of value, but many different explanations of its value are defended. George Agich notes this in his work on *Autonomy and Long-Term Care* where he describes both the philosophical and cultural perspectives on autonomy. In both cases he claims that autonomy is “generally, though not universally, regarded with approval.” [2] In support of Agich’s claim, there are, of course, the standard liberal arguments for the value of autonomy and these have been joined (with reservations) by communitarians, [3][4] feminists, [5][6][7][8] and bioethicists. [9][10][11] In these existing accounts, the focus has remained on either the formation and identification of

desires, beliefs, or characteristics or on the whole individual and not on the characteristics of human decision-making.[12] I will call these formal accounts of autonomy.

Formal accounts of autonomy share a relatively abstract conceptual focus, including formal requirements that describe the necessary *characteristics* or *capacities* of particular individuals or particular decision-making processes. There is a lack of corollary attentiveness to the *conditions* under which these decisions are made. Such an oversight limits the usefulness of an account for fostering autonomous decisions in health care (as well as other areas).

The Limits of Human Judgment

A common concern that motivates many conceptions of autonomy comes from *Brave-New-World*-like scenarios where an individual “believes” she is working in her own best interests, but in fact her actions are simply the result of social control. We do not want the thoroughly brain-washed individual’s decisions to count as autonomous. This demands certain conditions that eliminate obviously poor decisions. Examining how humans make decisions, however, shows that blatant poor decision making can arise much more subtly than through systematic deception. The mistaken assumption is that outside of relatively extreme conditions, individuals make good decisions, that they are at least instrumentally rational. That is, if an individual wants “P” and believes that “not-P unless Q”, then that individual will try to bring about “Q.” Research shows that people do not consistently (or even regularly) make instrumentally rational decisions. The obvious question is why not. In short, the answer is “bounded cognition.”

Suggested first by Simon, the bounded cognition framework assumes that individuals attempt to make optimal decisions, but that their decisions are affected by the conditions under which decisions are made. [13][14] These include but are not limited to: limits on available information, awareness of relevant criteria, constrained time, and limited memory. Importantly, because of the robust conclusions illustrated in the empirical research in this area, we are able to say more than that human decisions are biased—they are biased in predictable ways. [15]

One of the key concepts for this theory of human decision-making is the idea of a heuristic. A heuristic is a “short-cut” or “rule of thumb” in a decision-making process. For example, we operate under the condition of limited time. We do not have the opportunity to gather all the information that could be relevant to a decision we need to make. Instead of (literally) randomly choosing, we use a heuristic. This makes it possible for us to make decisions at a relatively rapid pace, but it also increases the likelihood that we will make mistakes. Heuristics are not bad or inefficient—in many cases they are much more efficient than laboriously researching every decision. Nonetheless, heuristics open up decision-making to certain likely mistakes.

There are at least three general categories of heuristics: representativeness, availability, and anchoring and adjustment. [16][17][18] Each heuristic opens up the possibility of several biases, but to make the case for effective autonomy, I will focus only on the heuristic of representativeness. In short, the representativeness heuristic arises from similarities in occurrences, outcomes and classifications. Although this heuristic is useful in many instances, it is also over-applied in many others leading to biases of judgment. For example, a “fair” coin will, on average, end up with a 50-50 split of heads

and tails. Individuals may expect a “fair” coin to flip heads, then tails, then heads, then tails, etc. In fact, a “fair” coin will not always (or even often) flip in this way.

Misapplying an objective characteristic (e.g., a “fair” coin) to a small data sample has come to be called the “insensitivity to sample size” bias arising from the representativeness heuristic.

This heuristic can also lead to particularly problematic judgments when it is used by a group of professionals. For example, there is evidence indicating that obese individuals are at greater risk for certain complications. [19] This evidence has been used by practitioners to justify excluding obese individuals from surgical interventions. Research by Everett et al., however, indicates that this claim is false in at least some cases. [20] Accordingly, deciding against surgery for obese patient in *all* cases is still an unjustified judgment. One explanation for the unjustified (but presumed justified) non-treatment of obese individuals by the medical profession is a biased use of the representativeness heuristic—specifically insensitivity to base rates. [21] As this example illustrates, failure to address the effects of bounded cognition can lead to poor decisions and, in health care, poor outcomes.

Procedure and Substance: An Uncertain Distinction

Formal accounts of autonomy have been separated into two camps: procedural and substantive accounts. A procedural account defines autonomy in terms of following certain procedures. If the appropriate procedures are followed, then the outcome is autonomous—regardless of the content of the outcome or the characteristics of the decision maker. Recent examples include procedural requirements like self-reflection,

freedom from inhibitions, identification or endorsement of desires, and at least minimal rationality. Substantive accounts differ from procedural accounts by limiting possible autonomous outcomes. They require certain conclusions or certain perspectives as a condition of autonomous choice. For example, Paul Benson includes the substantive requirement that an individual be able to competently criticize alternative courses of action according to the *relevant* norms. [22] This requirement is substantive because it restricts autonomy to persons with certain characteristics. No matter what procedures are followed in achieving the desire, an individual unaware of the relevant norms cannot be autonomous.

Substantive accounts of formal autonomy, typified by Paul Benson's account, are prone to devastatingly ambiguous normative standards. Benson initially claimed that for an individual to take a free action they must have "an ability to criticize courses of action competently by relevant normative standards". [22] Benson refined his claim such that an autonomous individual must view or understand him or herself as an agent competent to answer for their own conduct. Specifically, these answers would respond to "normative demands that, from one's point of view, others might appropriately apply to one's actions." [23]

Such a substantive account of autonomy is dogged with the indeterminacy of the "appropriate" normative standards. Are relevant norms those that an agent views as "right"? This makes the substantive account so weak as to be useless. If I get to pick and choose what norms to respond to, I will be able to competently respond for all of my actions by simply ignoring those norms that are the most difficult to respond to. Should it be those normative standards that are most widely accepted? It seems not, since this

would leave the determination of an individual as autonomous and the definition of relevance closely aligned with popularity. The dangers of such a perspective are illustrated in the familiar history of sexism, racism, and classism. Should it be those normative standards that are taken to be true independently of their popularity? Such a standard would require individuals to know the “true” from the “false” (or the relevant from the irrelevant) in order to qualify as autonomous. Additionally, this requires a method for determining the “true” from the “false” beyond currently available standards. And yet Benson never explains how we should determine relevance (or “truth”). Until such an explanation is offered, Benson’s account of autonomy will remain incomplete and so impractical.

There is another problem (or another way to get at the same problem) with Benson’s view. It is easy to imagine (or remember) an individual who has defied the norm of “reason-giving” to justify his or her actions. Think of the radical who describes the hegemonic views of bourgeois society as requiring reasons that are irrelevant or mistaken. Benson, of course, may defend his view (and accommodate this purportedly autonomous individual) by suggesting that some norms are invalid, and therefore, an individual need not be prepared to reply in accordance with them. But, again, this takes requires all autonomous individuals to know which norms are valid and which are invalid.

John Christman describes a formal conception of autonomy that I favor because it avoids the problems of substantive accounts as well as the problems of previous procedural conceptions of autonomy. [24][25][26][27][28] Following the literature, I refer to Christman’s account as a procedural account, [23][24][28][29] though,

momentarily, I will cast doubt on the distinction between substantive and procedural accounts. Christman identifies four conditions for the procedural autonomy of a desire. First, the formation of a certain desire was not (or would not have been) resisted when the individual did (or would have) reflected on the desire. Second, this lack of resistance is not due to circumstances that inhibit “self-reflection.” Third, the “self-reflection” is minimally rational and does not involve self-deception. Fourth, at a given time after the autonomous formation of a desire, the individual is minimally rational with respect to the desire. There are, of course, some ambiguities in Christman’s procedural conditions. Working definitions of self-deception and the conditions that inhibit reflection should be developed or imported. Moreover, fine-grained standards for manifest inconsistency and resistance (upon reflection) would be preferable. However, these ambiguities are unlikely to be devastating.

I favor something like Christman’s procedural account of autonomy over substantive accounts. However, Susan Wolf also defends a substantive account of autonomy that may be congruous with a procedural account. To be autonomous, Wolf claims an individual must be “able” to make decisions in line with “the True and the Good.” [30] Wolf emphasizes that what this entails is a commitment to some degree of “objective” evaluative standards. On her definition, even if no definitive conception of “the True and the Good” is achieved, an individual, to be autonomous, must be *able* to govern him or herself in accordance with standards independent of their own preferences. This has consistently been characterized as a substantive view of autonomy because it requires an agent to have an objectivist view of normative values. Even though no

specific values or aims are required, a certain perspective must be embraced for an agent to be considered autonomous.

It is worth asking whether Wolf's substantive view is very different from Christman's procedural view. Wolf requires the ability to govern actions according to "the True and the Good" and Christman requires reasoning that is without manifest inconsistency and self-deception. Both require outside judgment for determinations of autonomy. The self-deceived individual, by definition, cannot recognize their own limitations and the inconsistency of reasons implicitly acknowledges an objective standard. It certainly seems that Wolf's objectivism will require more than Christman's minimal rationality, but this may result from what we assume must be features of "the True and the Good." If we assume that this will include certain political or moral perspectives or certain well-defined commitments, then Wolf's account will go beyond what Christman requires. However, one might also read Christman's conditions as following from or even constituting the True and the Good (e.g., a life lived in accordance with the True and the Good is lived without desires and beliefs that have been or would have been resisted, without manifestly inconsistent reasoning and without self-deception). Whether or not these accounts can be happily combined, their potential similarity suggests that the distinction between substantive and procedural accounts of autonomy can be quite blurry.

With this in mind, my perspective on formal autonomy favors Christman's account but is not committed to a procedural account per se. Regardless, none of the accounts I have discussed (nor any other I can find) adequately address the heuristics and

biases of human judgment. I find this troubling as every potentially autonomous decisions we encounter is subject to human judgment.

Applying Bounded Cognition

As I noted before, bounded cognition marks the limits of formal accounts of autonomy. Even if we suppose that a desire is formally autonomous, because every decision is subject to the biases of human decision-making, some or many of an individual's choices will not reflect his or her autonomously formed desires. The desire may be formally autonomous, but the decisions may not reflect the desire.

We could, of course, ignore this disjunct and limit ourselves to protecting only formal autonomy. We would commit ourselves to value autonomy as a desire formation process and not in the actualization of the desire. On this view (assuming it can be adequately defended), achieving autonomously chosen ends would not be as significant as merely having (unattained) autonomous desires. Autonomously desiring to have a long life would be important, and actually having a long life would not (or would be less so). While I cannot deny the value of autonomously desiring a long life, I do think it is better to also have a long life. Tragically, some individuals will lack the means to have a long life, perhaps because of genetic or economic limitations. Worse, if we do not address the limits of decision making, these (and other) individuals will make predictably biased decisions that will further undermine their desire for long life.

Autonomy theorists may agree that effective autonomy is valuable, but argue that effective autonomy is irrelevant to the work that they are doing. In a certain sense, they are right. Defining formal autonomy is important whether or not formally autonomous

desires are actualized. However, in another sense, this view falls short of appropriate goals of responsible health care (and perhaps social policy in general). Even though focusing only on formal autonomy will be satisfactory for an abstract theory about autonomy, social practices and policies require more than formal autonomy. For example, if democracy is preferable based on the best social theories, we need to start by getting a grasp on the formal requirements of a democracy (one of which, coincidentally, is autonomous citizens). At the same time, we miss something important if we value democracy, but remain indifferent to whether or not social conditions (including voting mechanisms or overt coercion) undermine democratic election of public officials. In the same way, if autonomy is valuable, then we should work to define the formal features of autonomy (as so many have done). At the same time, we fail to take patient autonomy seriously and treat it, instead, as an esoteric and intellectual exercise if we remain indifferent to whether or not medical decisions reflect the formally autonomous desires of patients.

Effective autonomy clearly requires formal autonomy, but goes beyond it by requiring that decisions work toward the autonomously identified desires of the decision-maker. What is required for an adequate account of autonomy, an account that will foster effective autonomy is the following condition:

Effective autonomy is fostered just in case attempts are made to counteract predictable biases resulting from bounded cognition, including, when possible, the appropriate arrangement of background conditions and the use of responsible expertise.

Effective autonomy is fostered when biases resulting from bounded cognition are identified as risks in the decision-processes of individuals and this is followed by appropriate debiasing. What “appropriate debiasing” involves will depend on a number of variables for each decision including the likelihood of bias, the effect of bias (e.g., towards one alternative and/or away from another), the seriousness of the decision, and the cost of debiasing measures. The approximate value of these variables for the myriad of medical decisions requires empirical research, which has not, to date, been completed.

The Effect of Affect

The role of affect in demonstrating biased decision-making and in debiasing strategies remains unclear. For example, Salovey and Williams-Piehota did not pay methodological attention to the role of affect when they conducted field research that demonstrated how framing effects influence health-related decisions. [31] This shows that we can identify bias without identifying precisely what role (if any) affect plays in particular biases in decision-making. Modestly successful debiasing strategies have also been demonstrated in controlled and uncontrolled environments without attention to the role of affect. [32] [33]

Nonetheless, there is a growing body of evidence identifying the specific effects of affect on decision making. Research on affective forecasting (predicting one’s emotional response to a given event) has illustrated that mistakes in predicting affective responses (from the expected intensity and duration, to the expected affect itself) bias decision making in a number ways. [34] A study by Lerner and Keltner also

demonstrated that specific emotions can play a role in decision making. [35] Finally, Antonio Damasio postulates a specific role of affect in the availability bias regarding our views about safety in planes and cars. [36] Much like the effects of framing, the effects of affect (alongside predicted affect) warrant continued study as a set of conditions that can predictably bias decision-making. However, determining which debiasing strategies will require attention to affect requires more systematic evaluation of the role of affect in decision biases and corollary debiasing strategies.

In the next section, I suggest speculative policy changes that could arise from incorporating the heuristics and biases approach into a conception of autonomy for medical decision making. My suggestions center on the bias of overconfidence. Other policy changes will be needed in response to other biases, including those that arise from or are exacerbated by the affect of decision-makers.

Some Speculative Policy Changes Regarding the Bias of Overconfidence

It is difficult, and in some cases impossible, to evaluate whether a singular decision outside a controlled environment is biased. In controlled environments, the biases of a decision can be evaluated because the game can be fixed—we know what the accurate answer is. In the messy areas of health care (and many other decisions), the right answer for a single decision cannot always be clearly established in advance or even after the fact. However, there may be strategies that will help us avoid the predictable pitfalls of decision making. The debiasing strategies to pursue are not meant to foster “rationality” or eliminate “irrationality” per se, but to produce better outcomes by avoiding the predictable pitfalls in the decision-making process. A previously successful

strategy, recommended by Gigerenzer and Hoffrage, required replacing probability formats with frequency formats in describing the likelihood of events (e.g., the success of a treatment regimen or the accuracy of a diagnostic test). [37] This has been shown to improve the reasoning of naïve decision makers through the inclusion of this information in their decision process. In this section, I would like to suggest a couple of strategies that may improve clinical decision-making by avoiding the predictable pitfalls of overconfidence.

The bias of overconfidence—the systematic belief that we are right more often than we are—has been demonstrated repeatedly. [38] Individuals believe their judgments to be more accurate than they are. In cases where the accuracy of their prediction is important, this can be quite harmful—it may delay when they start chemotherapy, whether or not they take a medication, whether or not they consent to surgery. Of course, we cannot a priori decide that patients are mistaken about any particular judgment. Even if we have evidence that this particular patient is extremely overconfident, we remain unable to evaluate any particular judgment. This, however, does not leave us adrift. Studies by Hirt and Markman among others illustrated that considering-an-alternative debiased (to some extent) overconfident judgments about future events. [39][40] This effect, however, has also been shown to be limited by the ease of generating alternatives and the need for structure of the individual decision-makers. [41]

For medical practice, several possible policy implications arise. Whenever a patient makes a judgment about treatment, a routine prompt to consider multiple alternative outcomes may help calibrate patient judgments. For example, if a patient is considering chemotherapy and evaluating the side-effects, *routinely* prompting the patient

to consider the alternative (the side-effects will be more or less difficult/painful/exhausting) may produce better calibrated judgments. Whether or not some medical practitioners already do this, I am suggesting that this could be an expected routine, or “standard-of-care,” for every medical encounter.

Physicians as decision-makers are also susceptible to biases of overconfidence. Baumann et al provided some early evidence of overconfidence among medical practitioners. [42] Again, consider-the-alternative strategies may be successful here, but experts are in a special position to make similar decisions repeatedly. This is significant because confidence among experts can be better calibrated by consistent and clear feedback. [43] The need to provide clear and consistent feedback indicates several policy implications. Requiring an explanation for why another diagnosis is not accurate (and so an implicit consideration of alternatives) may be a useful routine. A systems-based policy requiring review by all involved physicians of every patient’s outcome may improve physician decision making even though it may be difficult and expense.

Though further research may show these recommendations to be unneeded, inefficient, or unsuccessful, we must recognize that a commitment to effective autonomy in medical practice requires attention to the role of predictable and correctable bias in decision making.

Conclusion

Effective autonomy is the coupling of decisions with formally autonomous interests—it requires formal autonomy, but goes beyond it by requiring attention to the background conditions in which decision makers will more accurately reflect their desires

through their decisions. Effective autonomy is fostered if and only if biases resulting from bounded cognition are identified as risks in the decision-processes of autonomous individuals and this is followed by appropriate debiasing.

The Corresponding Author has the right to grant on behalf of all authors and does grant on behalf of all authors, an exclusive licence (or non exclusive for government employees) on a worldwide basis to the BMJ Publishing Group Ltd to permit this article (if accepted) to be published in Journal of Medical Ethics and any other BMJPGJL products and sublicences such use and exploit all subsidiary rights, as set out in our licence (<http://JME.bmjournals.com/misc/ifora/licenceform.shtml>)

Competing Interests: None declared

References

1. McNeil B, Pauker S, Sox H, et al. On the elicitation of preferences for alternative therapies. *NEJM* 1982;306:1259-1262.
2. Agich G. *Autonomy and Long-Term Care*. New York: Oxford University Press 1993: 7, 14
3. Lehman E. ed. *Autonomy and Order*. Lanham, Maryland: Rowman and Littlefield 2000
4. Etzioni A. *The New Golden Rule: Community and Morality in a Democratic Society*. New York: Basic Books 1996
5. Friedman M. *Autonomy, Gender, Politics*. New York: Oxford University Press 2003
6. Hoagland S. *Lesbian Ethics: Toward New Value* Palo Alto, Calif.: Institute of Lesbian Studies 1988
7. Meyers D. *Self, Society and Personal Choice* New York: Columbia University Press 1989
8. Nedelsky J. Reconceiving autonomy: Sources, thoughts, and possibilities. *Yale Journal of Law and Feminism* 1989;1: 7-36
9. Schneider C. *Practice of Autonomy: Patients, Doctors, and Medical Decisions* New York: Oxford University Press 1998
10. Wolpe P. The Triumph of Autonomy in American Bioethics: A Sociological View in *Bioethics and Society*, eds. R. DeVries and J. Subedi. Upper Saddle River, New Jersey: Prentice Hall 1998
11. Katz J. *The Silent World of Doctor and Patient* New York: Free Press 1984
12. For example, take Berg JW, Appelbaum P, Lidz CW, and Parker. *Informed Consent: Legal Theory and Clinical Practice* New York: Oxford University Press 2001:vii. They provide (or attempt to provide) “a comprehensive introduction to the theory and practice of informed consent” for practicing clinicians, and yet, they never mention the psychology of decision-making.
13. Simon H. *Models of Man* New York: John Wiley & Sons 1957
14. March J and Simon H. 1958. *Organizations* New York: Wiley 1958
15. For a nice review, see Kahneman D, Slovic P, and Tversky A (Eds.). *Judgement under Uncertainty: Heuristics and Biases* Cambridge: Cambridge University Press 2001.
16. Bazerman M. *Judgment in Managerial Decision Making, 5th Ed.* New York: John Wiley & Sons 2002
17. Nisbett R and Ross L. *Human Inference: Strategies and Shortcomings of Social Judgment* Englewood Cliffs, New Jersey: Prentice Hall 1980
18. Tversky A, and Kahneman D. Judgment under uncertainty: Heuristics and biases. *Science* 1974;185: 1124-1131
19. Everett E, Tamimi H, Greer B, et al. The effect of body mass index on clinical/pathologic features, surgical morbidity, and outcome in patients with endometrial cancer. *Gynecologic Oncology* 2003;90: 150-157.
20. Hade A, Shine A, Kennedy N, et al. Both under-nutrition and obesity increase morbidity following liver transplantation. *Irish Medical Journal* 2003;96: 140-142

21. Base-rates refer to the standard frequency of a given event in a given context. So, for example, there is a base rate for college undergraduates who receive “A”s. Because this base rate has changed over the past several decades, a 4.0 grade point average means something different today than it did in the 1960s.
22. Benson P. Freedom and Value. *Journal of Philosophy* 1987;84: 465-486
23. Benson P. Free Agency and Self Worth. *Journal of Philosophy* 1994;91: 650-658
24. Christman J. Autonomy and Personal History. *Canadian Journal of Philosophy* 1991;21: 1-24
25. Frankfurt H. Freedom of the Will and Concept of a Person. *Journal of Philosophy* 1971;68: 829-839
26. Thalberg I. Hierarchical Analysis of Unfree Action. *Canadian Journal of Philosophy* 1978;8: 211-226
27. Dworkin G. *The Theory and Practice of Autonomy* Cambridge: Cambridge University Press 1988
28. Christman J. Constructing the Inner Citadel. *Ethics* 1988;99: 109-124
29. Christman J. Liberalism and Positive Freedom. *Ethics* 1991;101: 343-359
30. Wolf SR. *Freedom within Reason* New York: Oxford University Press 1990: 71
31. Salovey P and Williams-Piehota P. Field Experiments in Social Psychology: Message Framing and the Promotion of Health Protective Behaviors. *American Behavioral Scientists* 2004: 47 (5): 488-505.
32. Renner C and Renner M. But I Thought I Knew That: Using Confidence Estimation as a Debiasing Technique to Improve Classroom Performance. *Applied Cognitive Psychology* 2001; 15: 23-32.
33. Vaughn L and Weary G. Causal Uncertainty and Correction of Judgments. *Journal of Experimental Social Psychology* 2003; 39: 516-524.
34. Wilson TD and Gilbert DT. Affective Forecasting. *Advances in Experimental Social Psychology* 2003; 35: 345-411.
35. Lerner JS and Keltner D. Beyond Valence: Toward a Model of Emotion-Specific Influences on Judgement and Choice. *Cognition and Emotion* 2000; 14: 473-493.
36. Damasio A. *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: G.P. Putman's Sons 1994
37. Gigerenzer, G. & Hoffrage, U. How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review* 1995; 102: 684-704.
38. Fischhoff B, Slovic P, and Lichtenstein S. 1977. Knowing with Certainty: The Appropriateness of Extreme Confidence. *Journal of Experimental Psychology: Human Perception and Performance* 3: 552-564
39. Hirt ER and Markman KD. Multiple Explanation: A Consider-an-Alternative Strategy for Debiasing Judgments. *Journal of Personality and Social Psychology*. 1995; 69: 1069-1086
40. Koehler DJ. Explanation, imagination, and confidence in judgment. *Psychological Bulletin*. 1991; 110: 499-519.
41. Hirt ER, Kardes FR, and Markman KD. Activating a Mental Simulation Mind-Set Through Generation of Alternatives: Implications for Debiasing in Related and Unrelated Domains. *Journal of Experimental Social Psychology* 2004; 40: 374-383

42. Baumann A, Deber R, and Thompson G. Overconfidence among Physicians and Nurses: The 'Micro-Certainty, Macro-Uncertainty' Phenomenon *Social Science and Medicine* 1991; 32: 167-174.

43. Arkes HR, Christensen C, Lai C, and Blumer C. Two Methods of Reducing Overconfidence. *Organizational Behavior and Human Decision Processes* 1987; 39: 133-144.