The Virtues and Vices of Biased Rationality: An Eco-Cognitive Account

Emanuele Bardone (bardone@unipv.it)
Department of Philosophy, University of Pavia
P.zza Botta, 6 - 27100 Pavia, Italy

Lorenzo Magnani (lmagnani@unipv.it)
Department of Philosophy, University of Pavia
P.zza Botta, 6 - 27100 Pavia, Italy

Abstract

This paper aims at illustrating the virtues and vices of biased rationality. Starting with the virtues of biased rationality, we will illustrate the idea of *Homo Heuristicus* along with its fallacious dimension. We will claim that the rationale of biased rationality consists in turning ignorance into a cognitive virtue. More precisely, we will argue that biases provide us with premissory starting points even in absence of relevant information to solve the problem at hand. In the second part, we will turn to the vices of biased rationality: we will contend that the adaptive value of fallacies are limited to those situations in which we do not have (relevant) information, whereas *being* in such a situation is not adaptive at all. In the last part of the paper we will illustrate our main contention according to which debising rationality is made possible when human learning becomes an option ecologically supported by the accumulation of knowledge in the human cognitive niche. In order to provide ground to this idea, we will employ the lens model introduced by Egon Brunswik.

Keywords: biased rationality; fallacy; heuristics; de-biasing; lens model; cognitive niche

Introduction

This paper aims at illustrating the virtues and vices of biased rationality. Starting with the virtues of biased rationality, we will illustrate the idea of *Homo Heuristicus* along with its fallacious dimension. In this first section, we will claim that the rationale of biased rationality consists in turning ignorance into a cognitive virtue. More precisely, we will argue that biases (basically, resulting from fallacious reasoning) provide us with premissory starting points even in absence of relevant information to solve the problem at hand. This will permit us to develop the idea that a fallacy (or bias) does not necessary lead to a bad outcome. Accordingly, it can be either a good or bad line of argument. This is due to the fact that an argument is fallacious or biased with relation to a standard or a set of standards.

In the second part, we will introduce the distinction between *competence-independent information* and *competence-dependent information*. Building on this distinction, we will maintain that the adaptive role of biased rationality is conditional, as it lacks what we call *symptomaticity*. That is, the adaptive value of fallacies are limited to those situations in which we do not have (relevant) information, whereas *being* in such a situation is not adaptive at all.

The last part of the paper is an attempt to furnish a more general account about how decisions may be effectively unbiased by adopting a different perspective. Our main thesis is that, biased rationality turns out to be much less successful, as human learning becomes an option ecologically supported by the accumulation of knowledge in the human cognitive niche. In order to provide ground to this idea, we will employ the lens model introduced by Egon Brunswik to give ground to our proposal.

Introducing the *Homo Heuristicus*

In this section we are going to explore the very idea of biased rationality by illustrating the idea of *homo heuristicus* introduced and developed during the last two decades by Gigerenzer and colleagues (cf. (Gigerenzer & Selten, 2001; Todd & Gigerenzer, 2003; Gigerenzer, 2000; Gigerenzer & Brighton, 2009)). The idea of *homo heuristicus* explicitly addresses the problem of how to make two apparently conflicting concepts consistent: *accuracy* as the result of a certain decision, and *effort* as the amount of resources deployed in the decision-making process.

The idea of *homo heuristicus* stems from the rejection of two main assumptions about accuracy and effort. The first is that a heuristic always involves a trade-off to be reached between accuracy and effort, as they are basically conflicting concepts. In fact, accuracy usually involves time and resources. Therefore, given the fact that humans operate in cognitive economy with limited time and resources, they have to rely on decisions that are accurate enough, meaning that they might simply have to discard those strategies which lead to more accurate outcomes, but require greater resources. Heuristics are thought to be strategies reaching an accuracy-effort trade-off.

The second assumption can be called the "principle of total evidence". The principle of total evidence – introduced by Carnap (Carnap, 1947) and explicitly mentioned by Gigerenzer and colleagues – states that it is always better to take into account the total evidence available in order to determine whether or not a certain hypothesis or course of action is justified or rational: that is, having more information is always better than having less information. Or, to put it simply, *more is always more*, and *less is always less*.

Contrary to these two beliefs, Gigerenzer and colleagues argued, and managed to provide empirical evidence to support the idea, that heuristics are not always accuracy-effort trade-offs. On certain occasions, one can attain higher accuracy with less effort. Besides, more information may be detrimental leading not only to overload, but also to a general state of ignorance. Putting it simply, *less is more* and *more is*

less.

An example illustrating this point is the so-called "recognition-heuristic". What Gigerenzer and his team found is that when facing two alternatives, the one that is recognized is usually selected (Raab & Gigerenzer, 2005). In an interesting study, Raab and Gigerenzer asked two groups of university students respectively a German one and an American one, which city has a larger population between San Diego and San Antonio. Quite surprisingly, 100% of German students responded correctly, whereas only two thirds of American Students got the answer right. How could that be possible? We would expect American students to get it right, as San Diego and San Antonio are two American cities, and therefore they should know more about them or, at least, have more information. The explanation provided by Gigerenzer and colleagues is quite cunning. German students got it right, because they know less. More precisely, they got it right, because they only recognized one of the two cities, and thus they thought that it should have been the largest between the two. This is a fair example of recognition heuristic.

Gigerenzer and colleagues studied and tested a number of heuristics that turn out to be smart strategies for solving problems or making decisions. These heuristics compose what is called the "adaptive toolbox" (Gigerenzer, 2000; Gigerenzer & Brighton, 2009). Basically, this is a set of *fast-and-frugal* strategies that allows us to attain high accuracy while still operating in cognitive economy. For instance, *fluency heuristic*. In the fluency heuristic, both the of the two alternatives are recognized, but the one that is recognized faster is picked. So, going back to the example of the two cities, if I recognize San Diego faster than San Antonio, then I will choose it. Another example of smart heuristic is the so-called take the best. The best way to choose among concurrent and recognized options is to search for clues and stop as soon as one finds a discriminating clue favoring one above the others.

The idea of *less is more* is less paradoxical than one might think. It means that we simply operate selections: we basically search for the information that we think is relevant. Indeed, the selections we make are to some extent arbitrary, since we know what we select, but we do not know what we leave out. Heuristics are strategies that serve to this purpose: making use of what we know. In order to detail this point let us switch to the problem of the so-called fallacies. we will introduce and illustrate an alternative framework, which does not contrast with the results popularized by the proponents of the adaptive toolbox.

Easy to Use. The Rationale of Biased Rationality

The very idea behind the *homo heuristicus* furnishes an account about why biased minds make better inferences. In fact, heuristics are biases. Traditionally speaking, biases have been always considered as psychologically complex, leading to negative or unhappy outcomes. A bias is not necessarily an error, but it is usually considered as resulting from a poor or lower form of rationality, namely, biased rationality. They can indeed speed up a decision-making process, but, generally speaking, they are not necessarily a response to cognitive economy: they are easy to deploy. So, are they errors? Or not?

One can be biased, but at the same time achieve a good performance, just like in the cases illustrated by Gigerenzer and colleagues. For instance, a teacher may be biased towards certain students and those students might still perform very well at the same time – the so-called *primacy bias*. In this sense, it is not an error to get things right when judging on first impressions or coloring the interpretation of students' later performance based on initial results. It might however be an error because, if we got things wrong, we could be more easily blamed for our mistake and we would be told not to rush to judgment. We attribute a low cognitive status to biases because it is easy to highlight their weaker points, even though their weakest points do not necessarily lead us to be mistaken. It is this sort of *blamability* that is the source of mistrust. Basically, we have a commitment towards negotiating "the journey from cognitively virtuous starting points to cognitively virtuous outcomes" (Woods, 2009): that is, we start out safely and we want to arrive safely. Blamability warns us of the fact that we did not take the wrong path, but that it was a dangerous one.

More generally, we maintain that this attitude rests on the human capacity for planning ahead. Basically, when confronting a problem, people try to foresee possible objections, usually taking some precautions. Depending on their abilities and skills, people may anticipate some of the negative consequences a certain course of action might have. Planning ahead is somehow a certification that unhappy consequences may be prevented, even though we do not know precisely whether they are going to happen or not. Therefore, what we consider erroneous is the way biases (and fallacies) manage possible objections to a decision and/or unhappy consequences.

Appealing to Ignorance and Its Cognitive Virtue

Let us now go back to the question about which city has a larger population between San Diego and San Antonio. In order to clarify this point, we connect the recognition heuristic with an argument – traditionally deemed as fallacious – the so called *argumentum ad ignorantiam*. Let us make a very simple example. Suppose that John has to attend a meeting in the afternoon at his department, but he has not received any communication yet. Usually, department meetings are announced at least a few days before by the head of the department who sends an email to all the staff members. But this time she did not send any email to her colleagues. The meeting would usually start in less than one hour and John does not know what to do. Then, he carries out the following reasoning:

1. If there were a meeting at my department I would know it.

2. I do not know such a thing.

3. Thus there will not be a meeting today at my department.

This is considered by traditionally-minded fallacy theorists to be a fallacy. The main reason is that ignorance is never probative, meaning that it can only prove that one does not know a thing. In fact, in our example there might be a number of reasons why John did not come to know whether the meeting was going to happen or not. 1) Maybe his colleague sent him an email he never received, perhaps the head of department did not type his email address correctly, and then forgot to pay attention to the delivery failure message she should have received back. Carelessness in this case could be the reason explaining his ignorance. 2) Maybe his colleagues did not inform him about the meeting because they noticed he had not shown up during the last week, and they thought he had taken a week off to work on the last chapter of his book. Or 3) the head of the department and his other colleagues did not inform him on purpose, because they wanted to mob him.

In this case, we do not need to know whether John is eventually right or wrong but we can immediately see how it is easy – from an intellectual perspective – to raise some objections to John's argument. All the objections we were pointing to are related to the fact that John could have relied on a better argument.

John followed a pattern of reasoning that is labeled by AI theorists as *autoepistemic reasoning*. As Gabbay and Woods put it, "autoepistemic inferences are presumptive in character". Given that a candidate hypothesis is not known to be true, it is presumed to be untrue" (Gabbay & Woods, 2005). It is also known as negation as failure (Walton, 1995) or *argumentum ad ignorantiam*. An *ad ignorantiam* consists in an explicit appeal to our ignorance. In general, we analytically describe it as follows:

1. John knows he does not know P.

2. John asks himself whether he would have known P.

3. He would have known P, if it had been true.

4. He does not know P.

5. Then, he knows P is false.

What is interesting about this formulation is that it stresses how we are able to turn our ignorance into a cognitive virtue generating *premissory* starting points that we previously lacked. In fact, In 1) P is what prevents John from deciding. Conversely, in 4) P now becomes a clue suggesting a possible conclusion. It is 3) that describes the move allowing to escape ignorance without overcoming it. This ignorance-escaping feature of an *argumentum ad ignorantiam* should be treated along with another one: the move described in 3) is *ignorance-escaping* insofar as it is *irrelevance-avoiding*. In 1) John lacks premissory starting points that are relevant to the matter. In fact, in our brief debunking we brought up a number of objections that explicitly called for relevant information that should be acquired beforehand. However, although John does not overcome his ignorance by acquiring new and relevant facts, he escapes from it by avoiding any commitment to being relevant. That is, an *ad ignorantiam* does not get us out of ignorance, but it makes unapparent the distinction between relevant information and irrelevant information.

Relevancy avoidance is successfully performed, because it permits people to make a decision no matter what they know. The recognition heuristic chosen by the German group makes use of the same pattern of reasoning:

1. If I had known San Antonio, it would have been larger than San Diego.

2. But I do not know San Antonio.

3. Therefore, it is smaller than San Diego.

One important thing should be specified. The option generated by making use of our ignorance is not like the one we would gain by flipping a coin. It is certainly less arbitrary and more sophisticated, because it is at least a *spin*, as Woods argued. In fact, it permits us to make some guesswork possible. That is, we are not wholly in the dark. In fact, as already mentioned, an *ad ignorance* permits us to unfold premissory starting points, as they at least make a certain decision *decidable* or *affordable*.

Going back to our example, it is not true that John does not know anything. For instance, he knows something about what he should know. So, he can easily withdraw the hypothesis that he has not been informed because the head of the department wrongly typed his email address, because she always sets a return receipt option for such emails. She would let John know in the case she did not receive any confirmation from him. He could easily withdraw the second objection, because he usually works from home. And, as for the third one, he could discard that as well, because he has no problems at all with his colleagues.

The conclusion we are now arriving at is that an *ad ignorantiam* – belonging to biased rationality – is a weaker cognitive strategy than the one relying on relevant information.

The Vices of Biased Rationality

In the last section we made an explicit connection with the idea of *homo heuristicus* and his potentially fallacious dimension. In this section we will illustrate the problem of biased rationality going back to discuss fallacious reasoning with relation to the problem of relevancy. The treatment of this issue will be a crucial cornerstone in the introduction of our proposal.

In informal logic there is a class of fallacies identified as *ignoratio elenchi* (or *red herring*). By definition this class of fallacies introduces irrelevant information. The *argumentum ad ignorantiam* too can be classified as an *ignoratio elenchi*. Irrelevancy is always a relative matter, as it also depends on

the communicative context a certain reasoning is involved in (Bardone & Magnani, 2010). In this section we will try to make a step forward re-addressing the matter within a broader framework. This would also allow us to go beyond biased rationality.

Our main contention is that *ignoratio elenchi* is a kind of argument based on the introduction of what we call *competence-independent information*. Basically, an *ignoratio elenchi* is selected when a person does not have those competencies allowing him/her to address the original issue for debate. This kind of strategy may tell us something insightful about the nature of human decision-making and their cognitive system.

Consider for example a problem which many people face during their holiday: when to go swimming after eating. A doctor explains that there are several variables which we should account for in order to decide when to go swimming after eating. It depends on how much we eat, what we drink, the water temperature, whether we swim hard or not. All this information is *relevant* when deciding what to do and when to do it. Why is it relevant? It is relevant, because it would *explain* whether we may get cramps or other problems related to digestion. For instance, a heavy meal eaten just before swimming would make you feel sluggish and thus *explain* cramps.

As many people do not have the competencies a doctor has (or is supposed to have), they often rely on other kinds of information. For instance, mom's suggestions or what the majority do. By definition, mom's *sayso* or what the majority of people do are all irrelevant information. In our example, this information is irrelevant because it would not explain whether we may get cramps or not. More precisely, it is irrelevant because it is not *symptomatic*: what other people decide does not explain why a heavy meal affects our metabolism making us feel queasy.

The introduction and adoption of irrelevant information can be motivated by various reasoning. Indeed, it may be a strategy to divert audience attention and thus challenge the original issue for debate. Think for instance of how often politicians attack their opponents personally, not their ideas or the opinions they hold. However, as far as we are concerned here, we maintain that the introduction of irrelevant information is primarily a *cognitive strategy*, which responds to the necessity of *cognitive economy*. More precisely, it is a strategy that is deployed *in the absence of competence* regarding the matter in discussion. Thus, competence is connected with relevancy: being competent with regards to a certain matter is what permits employment of information which is relevant or, more precisely, symptomatic.

Focusing on competence so defined may help us solve or, at least, explore some open questions related to fallacy and biased rationality. First of all, what we argue is that the fallacious nature of biased rationality is concerned with the introduction of information that is not symptomatic to the conclusion that is drawn. Your mom's suggestions do not explain why you may have cramps or not; whereas what you have eaten does. The distinction between the two kinds of information can also be described in abductive terms: competence-dependent information are those which count as valuable clues guiding us to make the correct inference. In Peircean terms, such clues are those "from which we can infer that a given fact must have been seen" (Peirce, 1931/1958). In our view, irrelevancy is therefore an *epistemological* feature of fallacious reasoning and biased rationality and, above all, it characterizes them. Since the information we provide is not symptomatic, it is always *irrelevant*.

The question about competence and symptomaticity may also clarify the reason why a fallacy or bias is a *sometimes good, sometimes bad* strategy as pointed out by Woods (Woods, 2009). This was already clearly recognized by Aristotle who extensively argued around the unapparent defectiveness of fallacious reasoning. More precisely, the particular feature of fallacy is that it appears to have a certain property, when it has not. This unapparent defectiveness is connected with the fact that the information deployed in some fallacious reasoning is not symptomatic. As already pointed out, fallacious reasoning does not explain the reason why a certain event is *such and such*, and not *as such as such*. However, even if fallacies are not symptomatic, they can lead us to solve the task we are supposed to face.

The strategies based on competence-dependent information acquire an adaptive value, as they supply cognitive resources that are much more reliable than the ones based on competence-independent information. For instance, in the experiment about which city is larger between San Diego and San Antonio, it is most likely that an American expert in urban studies would not rely on recognition heuristics. More generally, experts tend to be *de-biased*, so to say. This is so, because knowledge is an increasingly reliable means for solving problems the more *abundant* it becomes. This is basically derived from the fact that knowledge is *resource-consuming*, whereas fallacies and biases are *resource-saving*. Strictly speaking, they are not corrigible, because they are optimized and ready to do their job. In this sense, arguments leaning on competence-independent information are neither corrigible nor enhancing.

This last contention is connected with a point introduced in the previous section, namely that fallacies and biases are somehow easy to dismiss. For instance, when people say that someone is biased or that an argument is fallacious they do not really mean that the person is mistaken or that the line of argument is wrong or false. They are merely pointing to a flaw in their opponent's reasoning, which might eventually lead to a bad decision or outcome, if followed.

So, being fallacious or biased renders an argument easy to dismiss. That one can be easily dismissed is not to be intended as a logical derivation. However weak or easy to dismiss an argument or bias is, it may allow a person to reach his target, as already maintained. In this sense, it describes a communicative move. For example, the simplest case of an *ad hominem* could help us make a decision which may

eventually be a good decision. Sometimes a person who has, for instance, a conflict of interest may be biased in holding certain positions. Therefore, knowing that he has a conflict of interest is not completely irrelevant insofar as it might be helpful, for instance, to my decision about how much weight to give to his claims.

More generally, our point is that irrelevance is a communicative feature. Irrelevance simply warns us that a certain piece of information may support an *easy-to-dismiss* point. Therefore, it prompts us to change or adjust our argument in order to acquire a better chance to succeed in a given discussion.

Appealing to Knowledge. De-Biasing Rationality

So far we have discussed the fact that biased rationality is indeed a survival strategy, however ill-grounded it may be from a more sophisticated perspective, namely, an intellectual one. The fact that we recognize some arguments or strategies as easy-to-dismiss simply means that we could have some better arguments and strategies at our disposal do we really have some better strategies or arguments? And what happens if we continue to use easy-to-dismiss strategies, even when they are not the only solution available?

As already pointed out, fallacies and heuristics make some decisions affordable for us, even though they do not resort to symptomatic information. This is due to the fact that biased rationality, as a set of fallacies and heuristics for decision-making, turns our ignorance into a cognitive virtue generating premissory starting points. So far so good. However, we do not overcome our ignorance, as we come up with premises that are not symptomatic or, at least, ambiguous. We overcome ignorance by establishing procedures that deliver information or resources that are somehow relevant or not biased. This can only happen by building up external structures that provide us with clues or information that are more symptomatic than the ones we previously had.

So, we introduce an eco-cognitive element, which will be the cornerstone of the argument we are now going to develop. We use the term eco-cognitive to stress that a behavioral option is available within a particular cognitive niche. Human beings owe their dominance over other animals to their ability to display advanced plastic behaviors. In turn, the possibility to display advanced plastic behaviors is closely related to having a second, non genetic, source of information, which upon occasion can deliver the proper resources to solve problems and help make decisions. Ultimately the ability to turn available raw materials into cognitive resources to support plastic responses is central to human success. The lens model theory introduced by Egon Brunswik (Brunswik, 1952, 1955, 1943) sheds some light on the dark side of biased rationality.

According to Brunswik, the relationship between the organism and the environment is defined by what he called "the lens model". The lens model is based on the idea that the relationship between the organism and the environment is mediated by the use of the so-called *proximal stimuli*, from which the organism can infer the distal state of the environment, which brought it about. *Ecological validity* is the term introduced by Brunswik to refer to the situation in which a given proximal stimulus acts as a valuable indicator of a certain distal state or event; ecological validity is a normative measure about *how diagnostic* certain proximal stimuli are with respect to a given distal event.

The main idea behind Brunswik's lens model is that it provides an alternative way to look into the questions related to domain-independent versus domain-specific approaches popularized by evolutionary psychology. His main contribution to this issue is to distinguish between the cognitive *process* of a certain activity and its *content*. He pointed out that the cognitive process of inferring a distal state of the environment from the proximal stimuli we received is *domain-independent*. Conversely, what is domain-dependent are those *indicators* or *local representatives* we make use of in order to infer distal states of the environment. For the indicator content is left unspecified.

In the light of Brunswik's lens model, adaptation (and thus the possibility of survival and reproduction) is the degree to which an organism attains a *stable relationship* with the external world (Kirlik, 2001, p. 238). In other words, achieving a stable relationship with the external world depends on developing *prepared* associations between a proximal stimulus and the corresponding distant event. Within this framework, plasticity is defined as that ability to change or adjust a prewired response according to the environment so as to increase the chance of a match between proximal stimuli and the distant state in an ever-changing environment.

As far as we are concerned here, plasticity can be defined as the ability to make use of those signs or clues that are more symptomatic of a certain event or situation than others. Ultimately, plasticity deals with the development of the abductive skills (Magnani, 2009), which allow us to detect clues and use them as indicators or local representatives of a distant event. In turn, these abductive skills basically rely on knowledge and competence. The crucial point for exploiting cognitive plasticity is to detect – and sometimes even create – various indicators specific to certain domains and not others in order to increase our chances of making successful inferences and judgments (Hammond & Steward, 2001).

This last contention leads us to consider cognitive strategies leaning on competence-independent information as ill-grounded for a long-term strategy insofar as it employs resources that by definition are not symptomatic. In Brunswikian terms, we may argue that the ecological validity of competence-independent information is quite poor, because it uses indicators that are not specific to a particular domain. Independence from a specific domain of application turns out to be the major limitation in this case. Mom's suggestions or following what the majority think are clues that cannot however be taken as reliable indicators or proximal stimuli of specific distal events.

In the previous discussion, we argued that some problems related to reproduction and survival are mandatory. One cannot cast them off, because that would impede reproduction and/or survival. Roughly speaking, under such conditions, giving an answer — even at random – is as good as giving the right answer. we posit that the virtue of strategies based on competence-independent information can only be *conditional*. That is, the *use* of competence-independent information in the situation where we have no information at all is "good", whereas *being* in that situation is not.

If so, one might expect that evolution would have provided human beings with a mechanism to escape from such conditions of having no information at all. Our contention is that such a mechanism is not provided at an individual level, but at the *eco-cognitive* one. The strategies based on competence-dependent information are adaptive as far as knowledge can persist and be accumulated and transmitted from generation to generation via the *cognitive niche*. As already pointed out, when knowledge is easily available, strategies based on competence may become dominant: easily available cognitive niches make abundance of knowledge possible.

More generally, individual agents spend part of their time tending to enhancement of cognitive assets if this makes the achievement of cognitive goals possible where they were previously unaffordable or unattainable (Magnani, 2007). Our claim is that this can only happen at the eco-cognitive level. Basically, a cognitive niche provides humans with an additional source of information storage and computational abilities, which support and even boost the capacity of exhibiting an increasingly flexible, adaptive response to an ever-changing environment (Magnani, 2009; Bardone, 2011). These extra-genetic materials, properly exploited by ontogenetic mechanisms like learning, provide the unique framework for re-adjusting and refining our cognitive assets *also* as individual agents.

Conclusion

In this paper we have discussed the virtues and vices of biased rationality. The main idea behind it is to draw a line distinguishing those situations when biased rationality can provide good solutions from those when it cannot, and therefore it should be dismissed. In pursuing that, we have introduced the distinction between *competence-independent information* and *competence-dependent information*. Building on such distinction, we have provided an alternative account of relevancy in decision-making by relying on the idea of symptomaticity. More generally, we contend that our contribution may serve two purposes worth mentioning here: 1) it helps clarify the concept of biased rationality for further empirical investigation; 2) it puts forward a theoretical framework potentially able to integrate psychological research with epistemological theory into a more coherent cognitive approach.

References

Bardone, E. (2011). *Seeking Chances. From Biased Rationality to Distributed Cognition*. Berlin/Heidelberg: Springer.

Bardone, E., & Magnani, L. (2010). The appeal of gossiping fallacies and its eco-logical roots. *Pragmatics & Cognition*, *18(2)*, 365–396.

Brunswik, E. (1943). Oranismic achievement and environmental probability. *Psychological Review*, *50*, 255–272.

Brunswik, E. (1952). *The Conceptual Framework of Psychology*. Chicago: University of Chicago Press.

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, *62*, 193–217.

Carnap, R. (1947). On the application of inductive logic. *Philosophy and Phenonmenlogical Research*, *8*, 133–148.

Gabbay, D., & Woods, J. (2005). *The Reach of Abduction*. Amsterdam: North-Holland. (Volume 2 of *A Practical Logic of Cognitive Systems*)

Gigerenzer, G. (2000). *Adaptive thinking: Rationality in the Real World*. Oxford: Oxford University Press.

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, *1*, 107–143.

Gigerenzer, G., & Selten, R. (2001). *Bounded Rationality: The Adaptive Toolbox*. Cambridge: Cambridge University Press.

Hammond, K. R., & Steward, T. R. (Eds.). (2001). *The Essential Brunswik. Beginnings, Explications, Applications*. Oxford/New York: Oxford University Press.

Kirlik, A. (2001). On Gibson's review of Brunswik. In K. R. Hammond & T. R. Steward (Eds.), *The Essential Brunswik. Beginnings, Explications, Applications* (pp. 238–242). Oxford/New York: Oxford University Press.

Magnani, L. (2007). Abduction and cognition in human and logical agents. In S. Artemov, H. Barringer, A. Garcez, L. Lamb, & J. Woods (Eds.), *We Will Show Them: Essays in Honour of Dov Gabbay* (pp. 225–258). London: College Publications. (vol. II.)

Magnani, L. (2009). *Abductive Cognition. The Epistemological and Eco-Cognitive Dimensions of Hypothetical Reasoning*. Berlin Heidelberg: Springer-Verlag.

Peirce, C. S. (1931/1958). *Collected Papers of Charles Sanders Peirce*. Cambridge, MA: Harvard University Press. (vols. 1-6, Hartshorne, C. and Weiss, P., eds.; vols. 7-8, Burks, A. W., ed.)

Raab, M., & Gigerenzer, G. (2005). Intelligence as smart heuristics. In R. J. Sternberg & J. E. Prets (Eds.), *Cognition and Intelligence. Identifying the Mechanisms of the Mind* (pp. 188–207). Cambridge, MA: Cambridge University Press.

Todd, P., & Gigerenzer, G. (2003). Bounding rationality to the world. *Journal of Economic Psychology*, *24*, 143–165.

Walton, D. N. (1995). *Arguments from Ignorance*. Penn State University Press: Philadelphia.

Woods, J. (2009). *Seductions and Shortcuts: Error in the Cognitive Economy*. (Forthcoming.)