2010 | 04

Working Paper

Research Department

Why do people give less weight to advice the further it is from their initial opinion?

By Francesco Ravazzolo and Øistein Røisland

NB NORGES BANK

Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles over e-post:
servicesenter@norges-bank.no
eller ved henvendelse til: Norges Bank, Abonnementsservice
Postboks 1179 Sentrum
0107 Oslo
Telefon 22 31 63 83, Telefaks 22 41 31 05

Fra 1999 og senere er publikasjonene tilgjengelige på www.norges-bank.no

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin e4ndelige form.
Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte.
Synspunkter og konklusjoner i arbeidene står for forfatternes regning.

Working papers from Norges Bank, from1992/1 to 2009/2 can be ordered by e-mail:
servicesenter@norges-bank.no
or from Norges Bank, Subscription service,
P.O.Box. 1179 Sentrum
N-0107Oslo, Norway.
Tel. +47 22 31 63 83, Fax. +47 22 41 31 05

Working papers from 1999 onwards are available on www.norges-bank.no

Norges Bank's working papers present research projects and reports (not usually in their final form)
and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested parties. Views and conclusions expressed in working papers are the responsibility of the authors alone.

ISSN 1502-8143 (online)
ISBN 978-82-7553-548-9 (online)

Why Do People Give Less Weight to Advice the Further it is from their Initial Opinion?*

Francesco Ravazzolo† and Øistein Røisland‡
Norges Bank

April 16, 2010

Abstract

Experimental studies on decision making based on advice received from others find that the weight put on the advice is negatively related to the distance between the advice and the decisionmaker's initial opinion. In this paper, we show that the distance effect can follow from rational signal extraction when the decisionmaker has imperfect knowledge about the advisor's competence. What drives the result is the assumption that the decisionmaker is better informed about her own competence than about the advisor's competence.

Keywords: distance effect, policy decision making, signal extraction, uncertainty.

JEL codes: C11, D78, D82, D83.

*We thank Beata Bierut and Krisztina Molnar for helpful discussions and for providing detailed comments. The views expressed in this paper are our own and do not necessarily reflect the views of Norges Bank.

†Norges Bank, Research Department. francesco.ravazzolo@norges-bank.no

‡*Corresponding author*: Øistein Røisland, Norges Bank, Monetary Policy Department. oistein.roisland@norges-bank.no

1 Introduction

Decisionmakers usually seek advice from other people. How decisionmakers take advice into account is an issue that has been subject to considerable research in organizational behavior, psychology, economics, and other areas. This paper considers one feature found in empirical studies about how decisionmakers use advice, which Yaniv (2004) labeled *distance effect*. The distance effect is characterized by a negative relationship between the weight placed on the received advice and the difference between the advice and the decisionmaker's initial opinion. The feature is documented by experiments in Yaniv (2004), Yaniv and Milyavsky (2007), and Onkal, Goodwin, Thomson, Gönül and Pollock (2009). In the psychology literature, the distance effect has been attributed to theories of attitude change (Aronson, Turner and Carlsmith (1963)), social judgment (Sherif and Hovland (1961)) and stereo type change (Kunda and Oleson (1997)). A common feature of all these explanations is that the distance effect violates a strict sense of "rationality".

In this paper, we offer an alternative explanation of the distance effect. Our approach models advice utilization as a Bayesian signal extraction problem. Generally, the optimal weight placed on the advice depends on the advisor's competence relative to the decisionmaker's competence. We assume that the decisionmaker is uncertain about the advisor's competence, as well as her own competence. The distance between the two agents' opinions can be used to update the estimates of the decisionmaker and the advisor competence. If the decisionmaker is more uncertain about her advisor's competence than about her own competence, the Bayesian updating scheme gives rise to the distance effect. This type of discounting of extreme advice is not related to traditional arguments for using various sorts of "trimming" in statistics, as considered by Yaniv (1997).

Our approach to analyzing the distance effect is similar in spirit to the approach taken by Benoît and Dubra (2007), who question the results from studies that claim to find evidence of that individuals tend to be overconfident about their abilities. Such studies typically find that a majority of people consider themselves better than the average (or median). Benoît and Dubra assume that agents follow a Bayesian updating scheme and show that such results are consistent with a rational use of limited information and can prevail even if the agents have an unbiased perception of their own competence. Even if we focus on a different "puzzle" than Benoît and Dubra (2007), our papers share the same assumption that an agent's apparent lack of rationality can be explained by her rational updating of estimates.

We stress that our result does not imply that people do not discount extreme advice due to the reasons proposed in the social-cognitive psychology literature. Our point that one cannot *exclude* the possibility that discounting advice according to its distance from one's initial opinion is a perfectly rational way for an agent to utilize new information given uncertainty about the advisor's competence.

2 The Model

Our model assumes that there are two agents; the decisionmaker (agent 1) and the advisor (agent 2). The decisionmaker challenge is to judge the value of an unknown variable μ. The decisionmaker has her own prior judgment (estimate) of the variables' value, and she receives a judgment from her advisor. Each agents' prior judgments are given by

$$x_1^{\texttt{prior}} = \mu + \epsilon_1 \quad (1)$$

$$x_2^{\texttt{prior}} = \mu + \epsilon_2, \quad (2)$$

where $E(\epsilon_1) = E(\epsilon_2) = 0$, $E(\epsilon_1\epsilon_2) = \sigma_{12}$, $E(\epsilon_1^2) = \sigma_1^2$, $E(\epsilon_2^2) = \sigma_2^2$. The expressions $x_1^{\texttt{prior}}$ and $x_2^{\texttt{prior}}$ can be thought of as (noisy) signals of the true value of the variable on which the decisionmaker shall form a judgment. We may interpret σ_1^2 and σ_2^2 as the (inverse of the) competence of agent 1 and 2 respectively. The decisionmaker combines the two signals according to

$$x_1^{\text{c,prior}} = (1-\alpha)x_1^{\texttt{prior}} + \alpha x_2^{\texttt{prior}} = x_1^{\texttt{prior}} + \alpha d, \tag{3}$$

where $d = x_2^{\texttt{prior}} - x_1^{\texttt{prior}}$.

Definition 1 *Let $\Delta \equiv |d|$ be the distance between the two agents' signals. The distance effect is a negative relationship between α and Δ.*

Known competence

First assume that the decisionmaker knows both σ_1^2 and σ_2^2. She uses both prior judgments to obtain a *combined* judgment. The optimal weight α placed on the advice (that is, minimizing the variance of the combined judgment error) is[1]

$$\alpha = \frac{\sigma_1^2 - \sigma_{12}}{\sigma_1^2 + \sigma_2^2 - 2\sigma_{12}}. \tag{4}$$

For simplicity, and with no consequences for the qualitative results, we assume that $\sigma_{12} = 0$, which gives optimal weight:

$$\alpha = \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2}. \tag{5}$$

Thus, the more competent the decisionmaker considers the advisor relative to herself, the larger weight she should place on the advisor's opinion. From equation (5), it follows that if the decisionmaker has perfect information about her own competence and the advisor's

[1] We assume that the advisor reports her judgment $x_2^{\texttt{prior}}$ honestly.

competence, the weight she puts on the advice is independent of the distance. Thus, the distance effect is not consistent with the optimal combined judgment under perfect knowledge about the respective competences.

Unknown competence

In practice, people rarely have full information about the competence they and others possess. As is common in the Bayesian approach, we assume that agent 1's prior expectations of her own competence and that of agent 2 follow an inverted gamma-2 density:

$$\sigma_i^{2,\texttt{prior}} \sim IG_2(\nu_i s_i^2, \nu_i), \quad i = 1, 2. \tag{6}$$

The degrees of freedom ν_i represent the prior belief of agent 1 on the mean value $\nu_i s_i^2/(\nu_i - 2)$: the higher is ν_i, the narrower is the prior density around s_i^2 and therefore the stronger is the prior belief[2]. The decisionmaker can still use the optimal combination scheme, but in a situation where α is based on her *estimates* of σ_1^2 and σ_2^2. Since it is only the *relative* competence of agent 1 and agent 2 that matters for the weight α placed on the advice, the crucial question is whether the realizations of $x_1^{\texttt{prior}}$ and $x_2^{\texttt{prior}}$ can be used for updating the judgments on the relative competence of both agents. Following previous assumptions, the variable d $(= x_2^{\texttt{prior}} - x_1^{\texttt{prior}})$ has mean 0 and variance $(\sigma_1^2 + \sigma_2^2)$. From (6) we have that the prior distribution on the sum $(\sigma_1^2 + \sigma_2^2)$ follows an inverted gamma-2 density, that is,

$$(\sigma_1^2 + \sigma_2^2)^{\texttt{prior}} \sim IG_2(\nu(s_1^2 + s_2^2), \nu), \tag{7}$$

[2]See, for example, Bauwens, Lubrano and Richard (1999) for moments of the inverted gamma-2 density.

where $\nu = \nu_1 + \nu_2$. Using the information from the distance $\Delta \equiv |d|$, the posterior estimate has the following closed form[3]:

$$(\sigma_1^2 + \sigma_2^2)^{\texttt{post}} \propto \nu(s_1^2 + s_2^2) + \Delta^2$$
$$= s_1^2 + s_2^2 + \theta(\Delta^2 - (s_1^2 + s_2^2)), \tag{8}$$

where $\theta = 1/\nu$. That is,

$$(\sigma_1^2 + \sigma_2^2)^{\texttt{post}} \sim IG_2(s_1^2 + s_2^2 + \theta(\Delta^2 - (s_1^2 + s_2^2)), \nu). \tag{9}$$

Thus, if the decisionmaker observes a large distance, meaning $\Delta^2 > (s_1^2 + s_2^2)$, she will judge the total competence of the two agents to be smaller (i.e., $\sigma_1^2 + \sigma_2^2$ higher). How can the decisionmaker use the distance Δ to update the estimate of the competence levels? It follows from equations (6) and (8) that agent 1 will update her estimate of the advisor's competence according to the following scheme:

$$\sigma_2^{2,\texttt{post}} \propto s_2^2 + \theta_2(\Delta^2 - (s_1^2 + s_2^2)), \tag{10}$$

where $\theta_2 = 1/\nu_2$. Similarly, agent 1 can update her estimate on her own competence according to

$$\sigma_1^{2,\texttt{post}} \propto s_1^2 + \theta_1(\Delta^2 - (s_1^2 + s_2^2)), \tag{11}$$

where $\theta_1 = 1/\nu_1$. Note that the decisionmaker will update the estimates of her own competence and the advisor's competence differently if $\nu_1 \neq \nu_2$, that is, when the confidence agent 1 has in her own competence is different than the confidence she has in the advisor's competence. Such a difference in confidence can be due to an *information asymmetry.*

[3] We follow standard Bayesian statistics and report only the kernel of the densities of interest. The sign "$\propto$" means "proportional" in our terminology and we use it to refer to the kernel of the distribution.

For example, it is reasonable to assume that individuals have better knowledge of their own competence than of other people's competence, in which case $\nu_1 > \nu_2$.

Agents can use posterior estimates in (10) and (11) to compute the posterior optimal weight α in (5) as:

$$\alpha \propto \frac{s_1^2 + \theta_1(\Delta^2 - (s_1^2 + s_2^2))}{(s_1^2 + \theta_1(\Delta^2 - (s_1^2 + s_2^2))) + (s_2^2 + \theta_2(\Delta^2 - (s_1^2 + s_2^2)))}. \tag{12}$$

The distribution for α is not a closed-form solution, but it can be derived by Monte Carlo simulations as the ratio between an inverted gamma-2 density in the numerator and the sum of two inverted gamma-2 densities in the denominator.

If the agents are not concerned with the estimation uncertainty, they will focus only on the mean weight $a \equiv E(\alpha)$; the value of a depends on parameters ν_1, ν_2, s_1^2, s_2^2, and the distance Δ. We do not focus on biases in the priors such as overconfidence, and therefore assume for simplicity that the decisionmaker has a prior expectation that she and her advisor are equally competent, meaning $s_1^2 = s_2^2$. Then, we have from (12):

$$a = \frac{1}{2} \quad \textit{with}\ \nu_1 = \nu_2. \tag{13}$$

If we assume an *information asymmetry*, for example $\nu_1 > \nu_2$, the distance effect emerges:

$$\begin{aligned} a &< \tfrac{1}{2} \quad \text{if } \Delta^2 > (s_1^2 + s_2^2), \\ a &= \tfrac{1}{2} \quad \text{if } \Delta^2 = (s_1^2 + s_2^2), \\ a &> \tfrac{1}{2} \quad \text{if } \Delta^2 < (s_1^2 + s_2^2). \end{aligned} \tag{14}$$

If agent 1 observes a large (relative to prior assumptions) distance Δ between her own signal and the advice (that is, $\Delta^2 > (s_1^2 + s_2^2)$), she discards agents 2' advice , reducing her weight a. If the distance is small (meaning, $\Delta^2 < (s_1^2 + s_2^2)$), agent 1 increases the weight on the advice of agent 2. The intuition for the result becomes clear if we assume that the decisionmaker has perfect information about her own competence, that is, $\theta_1 = 0$. Then, she will never update the estimate of her own competence, and the distance will only be used to update the advisor's competence. If she observes that the advisor has a judgment that is very far removed from her own judgment, she will (rationally) think that the advisor is less competent than she thought initially. Likewise, if the advisor gives counsel that is very close to her own judgment, she will increase her estimate of the advisor's competence and thus place a larger weight on the advice. The information embedded in the distance can be used to make a more precise estimate of the advisor's competence, and the distance effect is thus consistent with Bayesian updating under imperfect information about the quality of the advice.

2.1 Simulation examples

We analyze the relationship between the distance and the weight by two simulation exercises. In both exercises, we assume that the decisionmaker's prior assumptions of her own competence is equal to her prior assumption of the advisor's competence[4]. In Exercise I, she is equally confident about the estimate of her own competence as she is of her estimate of the advisor's competence, which implies that $\nu_1 = \nu_2$. In Exercise II, the decisionmaker is more confident about her estimate of her own competence than she is about her estimate of the advisor's competence, that is, $\nu_1 > \nu_2$.

[4]The analysis can be extended to consider estimation uncertainty, resulting in inferring the complete distributions for $\sigma_2^{2,post}$, $i = 1, 2$ and α. The parameter $\sigma_2^{2,post}$, $i = 1, 2$ will follow an inverted gamma-2 density, α can be computed by Monte Carlo simulations. For the sake of simplicity, we exclude this extension and focus only on mean results.

Table 1: Simulation design of exercises I-II

	EXERCISES	
PARAMETERS	I	II
σ_1^2	0.5	0.5
σ_2^2	0.5	0.5
ν_1	25	50
ν_2	25	5

Figure 1 illustrates how the decisionmaker utilizes the distance information for updating the estimates of her own competence and the competence of the advisor. If she observes a large distance Δ between her own signal and the agent 2's advice, the decisionmaker adjusts her estimates of their respective competence downwards (that is, increases the estimate of the variance of the signal error). Correspondingly, a small distance between the estimates signals that they are both more competent. In Exercise I, where the decisionmaker is equally confident of the priors (left panel), she adjusts the estimates of her own competence and her advisor's competence equally. In Exercise II (right panel), the decisionmaker is less confident about her prior estimate of the advisor's competence, and the distance has more informational value for estimating the advisor's competence than for estimating her own competence. The decisionmaker therefore revises the estimate of her advisor's competence more than the estimate of her own competence when observing the distance.

Figure 2 shows the relationship between the distance Δ and the weight a put on agent 2's advice, cf. equation (3). In Exercise I, observing the distance makes the decisionmaker adjust the estimates of her own competence and the advisor's competence equally, so that the weight put on the advice will remain unchanged. In Exercise II, the decisionmaker will adjust the estimate of the advisor's competence more than the estimate of her own competence, and the optimal weight put on the advice will therefore change. As seen from figure 2, the relationship between the optimal weight put on the advice depends negatively on the distance when $\nu_1 > \nu_2$. As argued above, it is reasonable to assume that people are more certain of their own competence than of other people's competence, which suggests that a negative relationship between Δ and a should be a natural tendency in practice. The *distance effect* is thus consistent with optimal signal extraction given a realistic information asymmetry.

Figure 1: Variances

The figures plot in the left panel prior values $s_1^2 = s_2^2$ (labeled "Prior") and posterior mean values $E(\sigma_1^{2,post}) = E(\sigma_2^{2,post})$ (labeled "1=2") for $\Delta = [0, 2]$ in exercise I. In the right panel prior values $s_1^2 = s_2^2$ (labeled "Prior"), and posterior mean values $E(\sigma_1^{2,post})$ (labeled "1"), and $E(\sigma_2^{2,post})$ (labeled "2") for $\Delta = [0, 2]$ in exercise II.

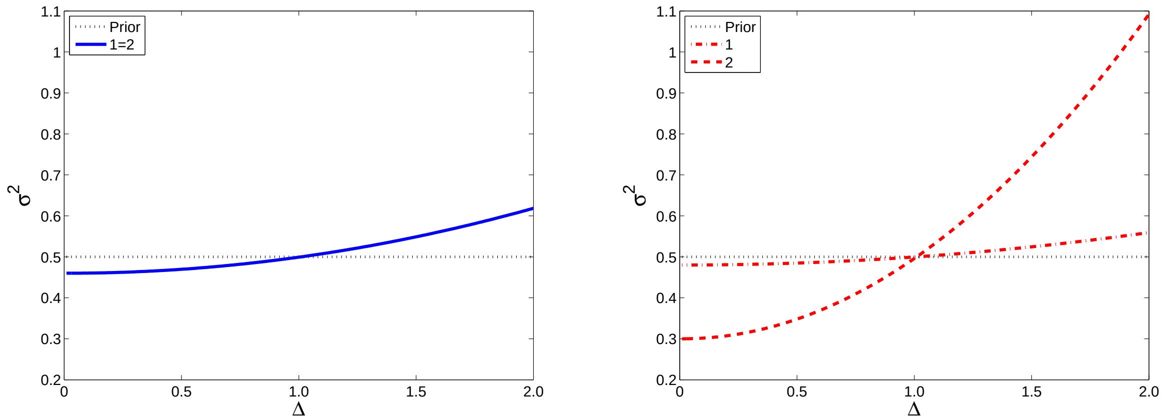

Figure 2: The relationship between Δ and a

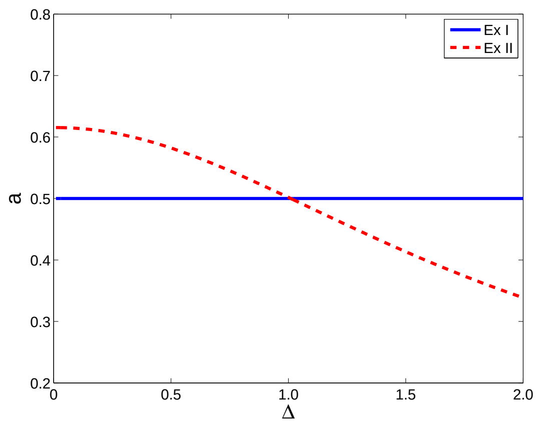

3 Discussion and Topics for Future Research

We have shown that the *distance effect* can be consistent with optimal signal extraction under asymmetric information. The key driving force is the assumption that the decision-maker is more confident about her prior assumption of her own competence than of her advisor's competence. So far we have not specified how these prior judgments are formed. One possible way to form the prior judgments on competency is to observe past signal (judgment) errors, i.e., ($\epsilon_i = x_1^{\texttt{prior}} - \mu$) from (1). This requires that historical values of the true variable μ are observable *ex post.* The estimate of the competence based on historical signal errors can then be computed according to[5]

$$\tilde{\sigma}^2_{t,i} = \frac{t-1}{t}\tilde{\sigma}^2_{t-1,i} + \frac{1}{1-t}(x^{\text{prior}}_{t,i} - \mu_t), \qquad i = 1, 2$$

where the time subscript t shows the dynamic structure of the updating.[6] If the decisionmaker has a longer series of her own past signal errors than of her advisor's signal errors, her own competence will be more precisely estimated than the advisor's competence, which would imply that $\nu_1 > \nu_2$. If the decisionmaker has had a long-term relation with the advisor, she has a longer series of signal errors than if the relation has been short. A testable implication of our model, which could discriminate between our explanation of the distance effect and the competing theories mentioned in the introduction, would be to investigate whether the distance effect is less prevalent the better the decisionmaker knows the advisor.

In some cases, our model would predict a reverse distance effect, that is, a positive

[5]See Weber (2008).

[6]Note than in the previous section, we focused on the signal extraction problem in a given period. The previous signal extraction problem was thus a static problem, and the time subscripts for the signals x_i^{prior} and the true value μ were therefore dropped.

relationship between the weight placed on advice and the distance from the decisionmaker initial judgment. For example, if the decisionmaker has little experience with the type of judgment she shall make, she will typically assume that she is not very competent, but she will also typically be very uncertain about her actual level of competence. If she receives advice from an experienced advisor, she could interpret a large distance between the advisor's judgment and her own judgment as an indication that she is even less competent than she thought, and conclude that she should therefore place larger weight on the advice. Hence, our model does not always imply that the weight given to the advice depends negatively on the distance from the decisionmaker's prior assumptions. The presence of asymmetric information implies that the weight placed on advice is not fixed, but depends on the actual distance from the initial assumption – the direction depends on whether the decisionmaker is more or less confident about her own competence compared to her advisor's competence. Nevertheless, we find it plausible that in most situations, people are better informed about their own competence than about other people's confidence, which suggests that the weight the decisionmaker places on the advice would tend to depend negatively on the distance from the prior expectation, as found in the experimental studies mentioned above.

In this paper, we have focused solely on the signal extraction problem of the decisionmaker and assumed that the advisor gives honest advice. However, in the principal agent literature, the focus is on agent's incentives to act strategically. An example is Prendergast's (1993) theory of "yes men", where she showed that it could be optimal for the principal to specify a contract that gives the agent's incentives to report judgments that are closer to their estimate of the principal's opinion (thereby meriting the term "yes men"). Our explanation of the distance effect gives, however, an alternative reason for the behavior “yes men”. To see this, assume that the advisor’s incentives are to maximize the weight that the decisionmaker puts on the advice. For example, the advisor’s salary could depend on the decisionmaker’s perception of the advisor’s competence. If there is a distance effect, meaning $\nu_1 > \nu_2$, and the decisionmaker thinks that the advisor is honest, the advisor would have an incentive to report a signal that is equal to the decisionmaker’s signal, since this would maximize the decisionmaker’s perception of the advisor’s competence. The existence of the distance effect therefore makes the case for “yes men”-like behavior stronger than it appears in the existing literature, since in Prendergast (1993), the agent would report signals that are biased toward the principal’s opinion but are not equal to the principal’s opinion. Arguably, it is not realistic to assume that the advisor acts strategically, while the decisionmaker thinks that that the advisor reports her signal honestly. An interesting topic for future research would be to study whether an equilibrium exists at all when the advisor utilizes the distance effect strategically and the decisionmaker knows the advisor’s incentives.

Another appealing extension would be to extend the one-shot game sketched above to a dynamic game with learning and reputational concerns as in Wrasai and Swank (2007). If the advisor’s future salary, or chance of reappointment, depends on the decisionmaker’s perception of the advisor’s competence, the advisor has an incentive to establish a “competent” reputation. One could then investigate the implications of the combination of a dynamic learning scheme and incentives for strategic behavior for the advice given the to decisionmaker and the weight that the decisionmaker places on the advice.

4 Conclusions

This paper focuses on the distance effect - a negative relationship between the weight placed on advice and the distance between the advisor's judgment and the decisionmaker's judgment. We show that the distance effect can be explained by using a Bayesian signal extraction model. If the decisionmaker is more confident about estimating her own competence than she is about gauging the advisor's competence, it follows from optimal signal extraction that a small distance between the decisionmaker's prior expectation and the advisor's recommendation would make the decisionmaker upwardly adjust her estimate of the advisor's relative competence, and thereby the weight put on the advice. Equivalently, a large distance would make the decisionmaker downgrade the advisor's relative competence. In our model, the existence of the distance effect hinges on an information asymmetry. We have argued that this type of information asymmetry is reasonable in practice, as people generally would be better informed about their own competence than about other people's competence.

One way to test our model would be to investigate through experiments whether the distance effect tends to be less prevalent the better the decisionmaker knows the advisor. The intuition for this test would be that when the decisionmaker has a relatively strong view of the advisor's competence, she would be less inclined to adjust this view when she sees the advice than if she has a very uncertain prior view on her advisor's competence.

References

Aronson, E., Turner, J. and Carlsmith, M.: 1963, Communicator credibility and communicator discrepancy as determinants of opinion change, *Journal of Abnormal and Social Psychology* **67**, 31–36.

Bauwens, L., Lubrano, M. and Richard, J.: 1999, *Bayesian Inference in Dynamic Econometrics Models*, Oxford University Press.

Benoît, J. and Dubra, J.: 2007, Overconfidence?, *Technical report*, Munich Personal RePEc Archive.

Kunda, Z. and Oleson, K. C.: 1997, When exceptions prove the rule: How extremity of deviance determines the impact of deviant examples on stereotypes, *Journal of Personality and Social Psychology* **72**, 965–979.

Onkal, D., Goodwin, P., Thomson, M., Gönül, S. and Pollock, A.: 2009, The relative influence of advice from human experts and statistical methods on forecast adjustments, *Journal of Behavioral Decision Making* **2(4)**, 390–409.

Prendergast, C.: 1993, A theory of "yes men", *American Economic Review* **83**, 757–770.

Sherif, M. and Hovland, C.: 1961, *Social judgment: Assimilation and contrast effects in communication and attitude change*, Yale University Press, New Haven.

Weber, A.: 2008, Communication, decision-making and the optimal degree of transparency of monetary policy committees, *Discussion Paper Series 1: Economic Studies 2008,02*, Deutsche Bundesbank, Research Centre.

Wrasai, P. and Swank, O. H.: 2007, Policy makers, advisors, and reputation, *Journal of Economic Behavior and Organization* **62**(04-037/1), 579–590.

Yaniv, I.: 1997, Weighting and trimming: Heuristics for aggregating judgments under uncertainty, *Organizational Behavior and Human Decision Processes* **69**, 237–249.

Yaniv, I.: 2004, Receiving other people's advice: Influence and benefit, *Organizational Behavior and Human Decision Processes* **93**, 1–13.

Yaniv, I. and Milyavsky, M.: 2007, Using advice from multiple sources to revise and improved judgment, *Organizational Behavior and Human Decision Processes* **103**, 104–120.