Information Processing and Commitment*

Armin Falk and Florian Zimmermann

June 27, 2017

Abstract

Beliefs are often found to be sticky and rather immune to new information. In this paper we highlight a specific mechanism that raises resistance to incorporate new information. We provide causal evidence that commitment to a first opinion leads to a neglect of new and challenging information. Being asked to write down an initial estimate before additional information is obtained significantly and substantially reduces the accuracy of final beliefs. Investigating the sources of this effect, we show that our findings are well explained by an internal desire to act consistently.

Many studies have documented that people frequently hold inaccurate and systematically biased beliefs. In the context of beliefs about own attributes, skills and abilities, individuals are often found to be overly optimistic (see, e.g., Svenson, 1981). The implications of such overconfident self-assessments have been studied in different contexts, ranging from tournament entry (Dohmen and Falk, 2011), CEO behaviour (Malmendier and Tate, 2008), overestimation of future gym attendance (DellaVigna and Malmendier, 2006) or excess entry into competitive markets (Camerer and Lovallo, 1999). Other studies have focused on beliefs about (external) states of the world. There, one of the most striking findings is the very high degree of belief divergence across individuals, even though the same sources of information are typically available. A prime example is climate change: Kahan (2012; 2013) demonstrates an enormous amount of heterogeneity and polarization in beliefs about whether climate change is human-made or not.

*Corresponding author: Florian Zimmermann, University of Zurich, Bluemlisalpstrasse 10, 8006 Zurich, Switzerland. Email: florian.zimmermann@econ.uzh.ch. We thank Steffen Altmann, Benjamin Enke, Andrea Galeotti, Simon Jäger, Matthias Wibral and two anonymous referees for many helpful comments that have greatly improved the paper. Financial support from the European Research Council is gratefully acknowledged. Parts of the results of this paper were previously circulated under the working paper "Preferences for Consistency".

This article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1111/ecoj.12542

These belief patterns arise even though individuals are often confronted with new information challenging their current views of themselves and the world, suggesting a pronounced immunity towards new (and possibly contradicting) information. This raises the question of which forces cause or contribute to the resistance to take novel information into account? In other words, which factors create the apparent immunity towards new information?

Taking these questions as point of departure, in this paper we highlight a specific channel that raises resistance to incorporate new information. Specifically, we report evidence from a series of experiments to make two key contributions. First, we provide clean and causal evidence that commitment to an opinion raises resistance to incorporate additional information. The resulting (partial) neglect of information causes lower accuracy of final beliefs and, in populations that are exposed to common signals, leads to higher belief heterogeneity. In a second step, we investigate the underlying mechanisms of the effect of commitment on information processing. The phenomena we identify is compatible with longstanding claims that people have a desire for behaving consistently (see, e.g., Cialdini, 1984). Making use of different treatment variations, we rule out potential alternative explanations and show that the commitment effect is caused by internal rather than social motives.

In our baseline experiments, subjects have to perform an estimation task and receive an additional signal regarding the solution of the task. In the main treatment, subjects are asked to state their current belief prior to receiving the signal and without knowing that they will later receive it. After they have received the signal they are free to revise their estimate. In the control condition, no prior statement is made. We find that the deviation between the final estimate and the signal is significantly larger in the main treatment compared to the control condition. Subjects in the main treatment take the new information to a lesser extent into account. Since all subjects obtain the same signal realization, as a consequence of commitment belief heterogeneity is higher in the main treatment in comparison to the control condition. In further treatment variations, we demonstrate the robustness of this finding and in addition show that commitment causes lower accuracy of final beliefs. In sum, our findings reveal that prior commitment to an opinion causes reluctance to incorporate new and challenging information. In our experiment, commitment is generated through the act of writing down a belief and stating it towards the experimenter. More generally, we interpret commitment as an action or statement that reflects the respective belief in an unambiguous way. Such commitment could be achieved through stating beliefs (about the self or some state of the world), for example towards friends, colleagues or family, but could also arise from actions that clearly reflect a certain belief, for instance buying a climate-friendly car.[1] As a consequence of commitment, existing biases in beliefs or more generally certain types of belief patterns persist, even in the presence of information challenging these patterns.

In the second part of the paper, we explore the underlying sources of the commitment effect. Our findings are compatible with a longstanding notion developed in social psychology, that people want to behave consistently. In the presence of such a desire, statements of beliefs or opinions become very important for subsequent information processing. Once individuals have committed to a certain belief, belief adjustments are only possible at the cost of being inconsistent.[2] Thus, even when being confronted with new and possibly challenging information, individuals stick to their prior stated belief in order to remain consistent.[3] In the second part of the paper, we first employ a series of

[1]The way we define commitment closely follows the social psychology literature on "consistency" (Cialdini, 1984). It is important to emphasise that this differs from the way the term is used in the economics literature. There, commitment typically entails an element of tying one's hands or irreversibility. In dynamic contexts full commitment implies the removal of all future choices. Instead, in our context, commitment by no means limits the choice set of decision-makers.

[2]In the Online Appendix we formulate a simple model that demonstrates how a taste for consistency can explain the pattern of findings from our experiment.

[3]The notion of consistency preferences also bears an interesting relationship to costs of lying. Gneezy (2005), Gneezy *et al.* (2013), Fischbacher and Föllmi-Heusi (2013), Abeler *et al.* (2014) examine lying behaviour in different contexts. They find evidence that subjects lie, but also that there is some cost of lying that prevents subjects from lying completely. A taste for consistency might contribute to lying costs, since lying might be viewed as a particularly drastic form of inconsistency. Most of the literature on consistency preferences has focused on the role of consistency as a means of social influence. Famous examples are Freedman and Fraser (1966) and Sherman (1980) who analyze the effectiveness

treatment variations to rule out two potential alternative explanations for our results, confirmatory bias and a combination of rational updating and differential effort provision. We then delve deeper and analyze if the commitment effect has internal or social roots.[4] Recall that in our experiment, commitment is achieved by writing down a first estimate and stating it towards the experimenter. Accordingly, both internal and social motives could be operating. We proceed in two steps: First, we show that purely internal motives are sufficient to generate the commitment effect. Second, in two different treatments we exogenously strengthen the social aspect of commitment. Both treatments do not further add to the commitment effect on information processing, suggesting that it is mostly driven by internal motives. In other words, the mere act of writing down an initial belief causes suboptimal information processing and as a consequence reduces belief accuracy.

Other explanations for why systematically wrong and divergent beliefs can be sustained have been suggested. In the context of overconfidence, one explanation that has been brought forward assumes a value of self-confidence and shows how overconfidence can be produced and sustained by selectively choosing information or by asymmetrically processing information, putting more weight on positive than on negative information (see for example Brunnermeier and Parker, 2005; Kőszegi, 2006; Falk *et al.*, 2006; Moebius *et al.*, 2013).[5] Another explanation stresses the role of social signaling for overconfidence (see Burks *et al.*, 2013; Charness *et al.*, 2014; Ewers and Zimmermann, 2015). If over-

of the so-called foot-in-the-door technique in different contexts, e.g., survey participation or engaging in embarrassing behaviour. The trick is to "tempt" a person to make a biased statement. In a second step he is confronted with a request related to that statement and the pressure to live up to it. Deutsch and Gerard (1955) show that commitment makes subjects less vulnerable to conformity bias.

[4] In the consistency literature, a central *internal* interpretation of a taste for consistency relies on the notion of cognitive dissonance. Pioneering work in this direction was developed in Heider (1946), Newcomb (1953) and Festinger (1957). Individuals want to behave consistently as this allows them to avoid cognitive dissonances. Relatedly, prior actions or statements might attract decision-makers' attention and focus, serving as an "anchor" and thereby affecting subsequent behaviour. A second interpretation instead emphasises *social* motives for consistency preferences. Falk and Zimmermann (forthcoming) provide a model where consistency preferences arise from a desire to signal skills towards others, building on an intuition in Cialdini (1984).

[5] Recently, there has been empirical support for biases in information-processing and information demand (see Eil and Rao, 2011; Moebius *et al.*, 2013; Charness *et al.*, 2014).

confidence is driven by a desire to signal skills or abilities to others, then individuals can appear overconfident even in the absence of any belief biases.[6] Concerning divergent beliefs about the state of the world (e.g., about the causes of poverty in society), two main classes of explanations have been brought forward. The first builds on the idea that information acquisition is costly (see, e.g., Piketty, 1995). At some point people simply stop acquiring new information and settle on the (incorrect) beliefs they have built so far. The second explanation is that beliefs are to a large extent shaped in a self-serving manner. People believe what they like to believe and interpret new information in that (self-serving) way (see, e.g., Bénabou and Tirole, 2006; Kahan, 2013). All these explanations seem important and jointly provide explanations for how belief biases and belief divergence can be sustained. Our study highlights the role prior statements of beliefs play in raising resistance towards new information. Note that none of the theories discussed above can account for our findings. Two further points concerning our mechanism seem noteworthy and distinguish it from existing explanations: First, it does not depend on the context, i.e., it can explain stickiness in beliefs about the self as well as beliefs about external states of the world. Second, it applies irrespective of the process that generates belief patterns. Whatever the underlying process that formed existing beliefs, commitment is likely to create persistence.

The remainder of the paper is organised as follows: In the next section we establish the commitment effect on information processing. In section 2 we investigate the sources of the commitment effect, and section 3 concludes.

[6]Notice that Benoît and Dubra (2011) have provided a different explanation for overconfidence and argue that in fact most of the evidence for relative overconfidence can be reconciled by correct Bayesian updating from common priors. In other words, evidence in the form "40% of subjects place themselves in the top 20% of good car drivers" should not be interpreted as evidence for overconfident self-assessments as it can be the outcome of correct updating from unbiased information. Interestingly, Benoît *et al.* (2015) report results from two experiments that cannot be produced by rational Bayesian updaters, providing clean evidence for overconfidence.

1 The Effect of Commitment on Information Processing

The goal of our experiment was to identify the role of commitment for information processing in a clean and unambiguous way. We implement an estimation task environment where subjects have some (private) information and then obtain an additional piece of information about the solution of the task. The key idea of our experimental design is to create treatment conditions that are completely identical, except that in some conditions subjects will be asked to commit to a belief before an additional signal is received.

We conducted a total of seven treatments to identify the commitment effect on information processing as summarised in *Table 1.* The table also summarises experiments that were conducted to obtain the additional piece of information as well as additional treatments that allow us to uncover the sources of the commitment effect (these treatments will be discussed in detail in section 3).

1.1 Design

Our key treatment comparison is always between one condition with commitment and one without. The different steps of the experiments are illustrated in *Figure 1.* We start by describing our two principal treatments, **Commitment** and **NoCommitment**. The **Commitment** treatment is shown in the upper panel. First, subjects were explained the task: They had to estimate the number of peas in a bowl.[7] Subjects were paid according to the precision of their estimate. If their estimate was less than 5 percentage points above or below the true value of 3000, subjects earned 10 euros. For every 5 percentage points the estimate deviated from the true value, we deducted 50 cents. For example, a subject whose estimate deviated 17 percent from the true value earned 8.50 euros. Negative earnings from the estimation task were not possible.

[7] A picture of the bowl is shown in Online Appendix A.

Subjects were seated around a table, which was placed in the middle of the lab.[8] After subjects had been informed about the task the bowl was shown. The bowl with peas was placed in the middle of the table and subjects had time to think about an estimate, i.e., to form a belief about the correct number of peas. Subjects were asked to raise their hand once they had written down their estimate on an answer sheet that had been distributed at the beginning of the experiment. As soon as a subject indicated that he had written down an estimate, the experimenter recorded the subject's estimate. This means that subjects had written down their first estimate and knew that the experimenter knew that estimate. At this point, subjects had committed to their first estimate.

After all subjects had stated their estimates, the experimenter announced that he would now provide subjects with additional and "helpful" information regarding the estimation task.[9] Each subject received an information sheet containing the following sentence. "In the past it has often been the case in various estimation tasks, that the average estimate of all participants is often relatively close to the true value. The estimation task you are facing has also been conducted with a different group of participants. They have also been paid according to precision of their estimates. The average estimate of the number of peas in the bowl of this group was 2615. If you want to, you can now revise your estimate." After they received the information sheet, subjects had time to revise their estimate on their answer sheet. Of course, only the final estimate was relevant for earnings. After all subjects had indicated that they had specified their final estimate, answer sheets were collected and the estimation task ended.

The additional information we provided to subjects was based on a separate experiment we had conducted with 61 different subjects (see *Table 1*). They faced the same estimation task and were also paid according to the precision of their estimates. The average estimate of that group was 2615. Note that the additional information was the

[8]Subjects were seated sufficiently far away from each other, such that they could not see what other subjects were writing down.

[9]Note that in both treatments the additional information came as a surprise, i.e., subjects did not know that they would receive the additional signal before they actually received it.

Figure 1: *Timing of the Experiment*

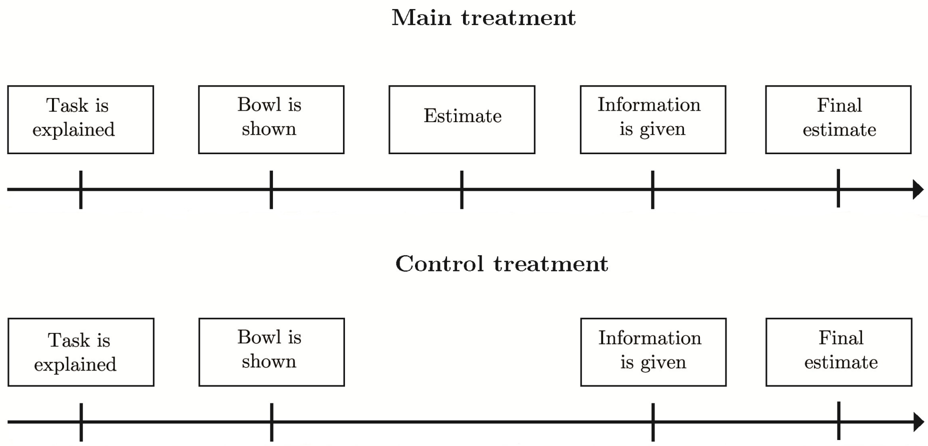

same for all subjects, i.e., it was a public signal available to all subjects.

Subjects in the **NoCommitment** treatment also learned about the task, i.e., that they needed to provide an estimate about the number of peas in the bowl. The only difference between the two treatments was that in the latter subjects did not state an estimate prior to receiving the additional informative signal (see lower panel of *Figure 1*). Subjects saw the bowl with peas for some time prior to receiving the information sheets. The time was approximated to be the same as the time subjects needed in the Commitment treatment. During this time subjects could form a belief about the correct number of peas, but did not write it down or state it to the experimenter, i.e., no commitment to a first estimate was made. After subjects received the sheet containing the additional information, they stated their estimate on their answer sheet.

Treatments Commitment and NoCommitment allow us to investigate whether commitment to a first estimate affects to what extent subjects incorporate additional information in their final estimate. But we are of course also interested in the effects of commitment on the actual *quality* of information processing. In our opening examples, we described situations where people appear to be rather immune to new information that challenges their current views of the world. Thus, the question arises whether we can

Table 1: *Experimental Treatments, Descriptions and Summary Statistics*

Treatment	Description	# Subjects	Dates	First Estimate Mean (s.d.)	Absolute Deviation from Signal Mean (s.d.)	Group Size Indifference Mean (s.d.)
Commitment	common signal of 2615	54	Winter 2010	2373.2 (2521.8)	912.1 (796.5)	NA
NoCommitment	common signal of 2615	51	Winter 2010	NA	570.4 (614.2)	NA
Commitment10	different signals, based on groups of 10	50	Fall 2015	2447.8 (2365.0)	1136.6 (1213.9)	NA
NoCommitment10	different signals, based on groups of 10	50	Fall 2015	NA	590.4 (677.5)	NA
CommitmentNormal	different signals, drawn from normal distr.	47	Fall 2015	2504.8 (2426.1)	750.6 (820.6)	NA
NoCommitmentNormal	different signals, drawn from normal distr.	47	Fall 2015	NA	413.1 (425.3)	NA
BeliefAccuracy	valuation of own versus group precision	49	Fall 2015	NA	NA	4.0 (2.2)
Group of 61	signal in (No)Commitment was based on	61	Winter 2010	2615.1 (2073.2)	NA	NA
Group of 200	signals in (No)Commitment10 were based on	200	Winter 2010/ Fall 2015	2718.7 (2478.0)	NA	NA
CommitmentR	robustness check with 5 min sharp for task	50	Fall 2015	2438.4 (2487.1)	862.0 (843.9)	NA
NoCommitmentR	robustness check with 5 min sharp for task	49	Fall 2015	NA	548.7 (498.9)	NA
NoCommitmentOwnHyp	ensured that subjects formed own hypothesis	47	Spring 2012	NA	588.1 (533.3)	NA
CommitmentNoSI	removed the social comp. of commitment	54	Spring 2012	NA	964.6 (703.4)	NA
CommitmentSI	observed by bystander	50	Fall 2015	2306.9 (1845.3)	906.8 (699.36)	NA
CommitmentSImoney	observed by bystander and 3 Euro if selected	48	Fall 2015	2273.4 (2261.3)	1013.0 (821.6)	NA

Notes: Short descriptions and summary statistics for all our experimental conditions. *Group Size Indifference* captures the outcome variable of the Belief Accuracy treatment, i.e., the group size at which subjects started to prefer being paid according to the group average rather than their own own estimate.

find evidence that commitment to an opinion raises resistance to incorporate new information, i.e., reduces the quality of information processing? Since we are not operating in a Bayesian framework, we cannot use proximity to the Bayesian posterior as a measure of quality. Instead, based on the idea that making better use of the additional information should on average (and everything else equal) lead to more accurate final estimates, we take the accuracy of final estimates as a measure of processing quality. However, our two treatments from above cannot be used in this respect, because in these treatments, all subjects received the exact same piece of information. Thus, accuracy of final estimates reflects the accuracy of the one specific signal realization we provided to subjects, rather than the optimality of information processing as such. Therefore, we conducted additional treatments where different subjects obtained different signal realizations, rendering estimation accuracy a meaningful measure of processing quality.

Treatments Commitment10, NoCommitment10, CommitmentNormal and NoCommitmentNormal were identical to Commitment and NoCommitment respectively, except for the additional information we provided. For **Commitment10** and **NoCommitment10** we first had a total of 200 different subjects provide an estimate. We randomly divided the 200 subjects into subgroups of 10, and used the average estimates of the subgroups as pieces of information, thereby obtaining 20 independent pieces of information.[10] Each subject in **Commitment10** was randomly assigned to one such piece of information. **NoCommitment10** served as control condition, and we ensured that always pairs of subjects (one subject from Commitment10 and one subject from NoCommitment10) obtained the same piece of information. In addition, we were interested in how subjects assess their own estimation ability relative to the average estimate of a group of subjects. Specifically, we wanted to know for which group size subjects would prefer to be paid according to the average estimate of the group rather than their own estimate. In treatment **BeliefAccuracy**, we employed a price list format where subjects, after seeing the bowl with peas for the same time as in the NoCommitment treatments, in a series of choices could always choose to be paid according to their own estimate or a group average, where the group size was varied in increments of 1 from 1 to 10.[11] This provides

[10] The way the piece of information was introduced was identical to (No)Commitment, except for the last sentence where subjects were informed that the additional information was based on the average estimate of a group of 10 subjects (see also Online Appendix C). Actual signal realizations ranged from 2069 to 3307. See also *Table 1* for summary statistics on the group of 200.

[11] Notice that subjects did not know the actual average estimates of the respective groups (i.e., the signal realizations) when making their choices. They only knew that if they chose the respective option, they would be paid according to the average estimate of a group of the respective size. After filling out the price list, subjects were asked to provide an own estimate, in case this would be their payment-relevant estimate.

us with a measure of subjects' relative valuation of the signal (see Online Appendix C).

In **CommitmentNormal** and **NoCommitmentNormal** we moved one step further by making the additional information more objective. Instead of providing subjects with the average estimate of a group of subjects, we provided them with a random draw from a normal distribution with mean equal to the true state of the world (3000) and a standard deviation of 300. The signal distribution was explained to subjects in great detail. Subjects were provided with a graphical representation of the normal distribution, and the unbiasedness of the signals was emphasised (see Online Appendix C)). Each subject in **CommitmentNormal** obtained an independent random draw from the normal distribution.[12] **NoCommitmentNormal** served as control condition, where we again made sure that always pairs of subjects (one subject from CommitmentNormal and one subject from NoCommitmentNormal) received the same signal realization.

1.2 Procedural Details

A total 348 subjects participated in the seven treatments (see *Table 1*). Subjects were mostly undergraduate students from various fields at the University of Bonn and were recruited using the online recruitment system by Greiner (2003). No subject participated in more than one session. The experiments were conducted with paper and pencil. Instructions, translated into English, are provided in Online Appendix C. Sessions lasted on average about 40 minutes and subjects earned about 12 euros on average, including a show-up fee of 5 euros.

1.3 Hypotheses

In all commitment / no commitment treatments, subjects see the bowl with peas and form a belief about the correct number of peas. Then subjects receive an additional piece of information. The standard prediction would be that the way subjects process this

[12] Signal realizations ranged from 2235 to 3748 (mean = 2966.7, standard deviation = 288.32).

information does not depend on commitment, and accordingly the (absolute) deviation between final estimate and signal and between final estimate and the true value of 3000 should not depend on commitment. As a consequence neither belief heterogeneity nor belief accuracy should depend on prior commitment.

If, however, prior commitment to a first estimate raises resistance to take novel information into account, we should see that subjects in the treatments with commitment take the information to a lesser extent into account. As a consequence, the absolute difference between final estimates and the signal should be higher in the treatments with commitment compared to those without. For treatments Commitment and NoCommitment, since all subjects obtain the same signal realization, commitment should also lead to higher belief heterogeneity, compared to the NoCommitment condition. For treatments Commitment10, NoCommitment10, CommitmentNormal and NoCommitmentNormal, final estimates should be less accurate in conditions with prior commitment. In Online Appendix B we show formally that this set of predictions can be generated from a model where individuals have a taste for consistency.

HYPOTHESES: *The absolute difference between final estimates and the informative signal is higher in the treatments with commitment relative to the treatments without commitment. Belief heterogeneity will be higher in Commitment, relative to NoCommitment. Final estimates in Commitment10 and CommitmentNormal will be less accurate relative to NoCommitment10 and NoCommitmentNormal, respectively.*

1.4 Results

We start by analyzing findings from our two main treatments, **Commitment** and **NoCommitment**.[13] Pooling data from both treatments, the average final estimate was 2411.44. The estimation problem is rather difficult for subjects and final answers ranged from 400 to 6000. Accordingly, the variance was large as indicated by a standard de-

[13]The online material contains the dataset as well as a do-file that reproduces the analysis.

viation of 1026.81. We chose a difficult task on purpose as it offers an ideal context to provide subjects with an additional signal and study subsequent information processing.

We now turn to our main variable of interest, the absolute deviation between final estimates and the signal of 2615. In *Figure 2* (upper left panel) we show a histogram of absolute deviations. The figure reveals that extreme deviations from the signal are more frequent in the Commitment treatment than in the NoCommitment treatment. In contrast, small deviations are relatively more frequent in the NoCommitment treatment. For instance, about 47 percent of all estimates in the NoCommitment treatment are in the interval +/-300 around the additional signal, compared to only 26 percent in the Commitment treatment. On average, the deviation in the Commitment treatment is 341.7 points higher than in the NoCommitment treatment. The difference in deviations from the additional signal is statistically significant, as shown in *Table 2* (column 1) where we regress the absolute deviation from the signal on a constant and a treatment dummy.[14] In other words, estimates in the main treatment deviate more strongly from the signal of 2615, documenting that commitment indeed induced subjects to incorporate the new piece of information to a lesser extent.

RESULT 1: *Statements of a first opinion induced subjects to incorporate the new piece of information to a lesser extent. The absolute deviation between final estimates and the public signal is higher in the Commitment treatment, compared to the NoCommitment treatment.*

Next, we look at belief heterogeneity in the two treatments. Recall that in Commitment and NoCommitment, the additional information was the same for all subjects, and hence in principle should reduce belief heterogeneity by the same amount. In light of our previous finding, however, we would expect to see higher belief heterogeneity in the Commitment treatment. This is indeed what we find. The standard deviation of final

[14] All findings from our treatment comparisons are robust to using the non-parametric Wilcoxon Ranksum test instead of linear regressions.

Figure 2: *Histograms of Absolute Deviations*

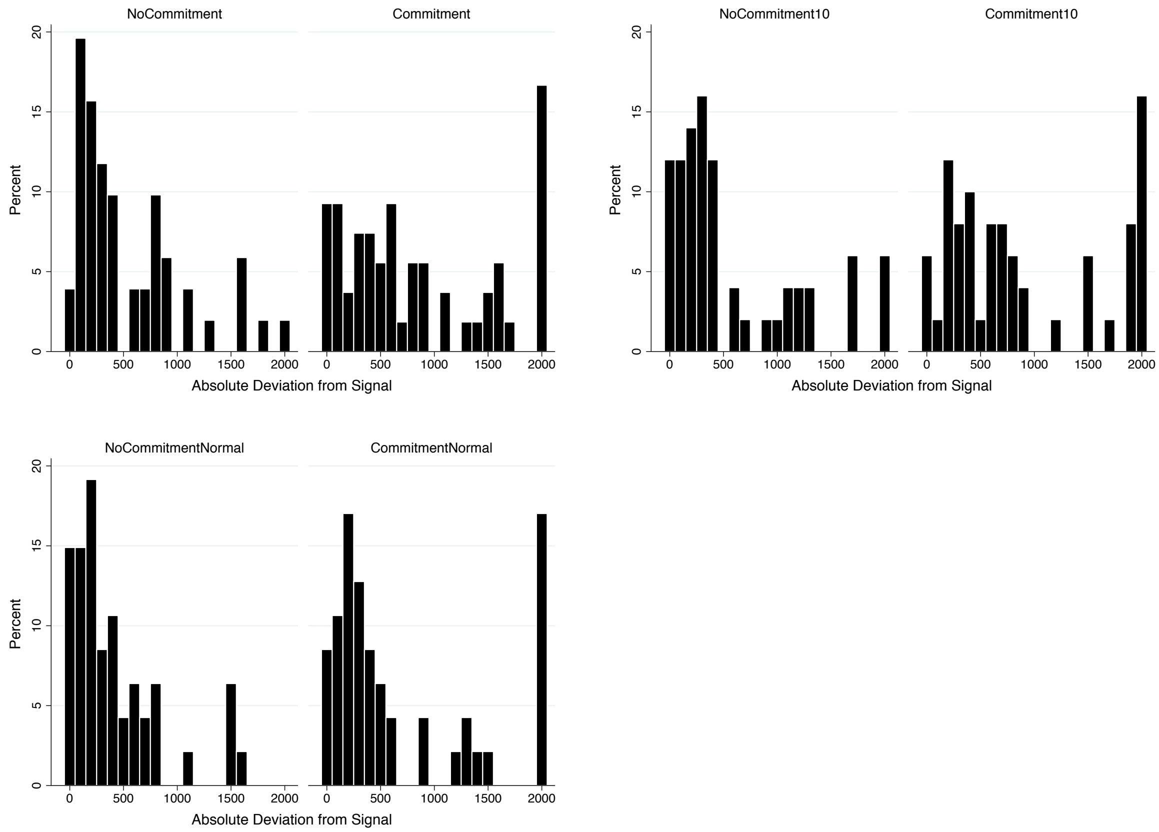

Notes: Histograms of absolute deviations from the signal. Observations are censored at 2000. *Upper left panel* shows treatments NoCommitment and Commitment, *upper right panel* depicts treatments NoCommitment10 and Commitment10, and *lower panel* shows NoCommitmentNormal and CommitmentNormal.

estimates is 1185.95 in the main treatment, compared to 831.68 in the control condition. This difference is significant, using Levene's test for differences in variances (p-value = 0.01).

RESULT 2: *Belief heterogeneity is higher in treatment Commitment, compared to treatment NoCommitment.*

In a final step we investigate whether commitment is associated with a decrease in the quality of information processing, i.e., the quality of final estimates. So far we demonstrated that commitment to a first estimate reduces the extent to which subjects incorporate additional information in their final estimate. But is there also evidence that the quality of information processing decreases through commitment? We conducted treatments Commitment10, NoCommitment10, CommitmentNormal and NoCommitmentNormal to address this question.

Table 2: *Linear Regressions*

	(1)	(2)	(3)	(4)	(5)
	Commitment and NoCommitment	Commitment10 and NoCommitment10		CommitmentNormal and NoCommitmentNormal	
	Dependent Variable				
	Deviation Signal	*Deviation Signal*	*Inaccuracy*	*Deviation Signal*	*Inaccuracy*
Commitment	341.7** (138.4)				
Commitment10		546.2*** (196.6)	537.6*** (182.7)		
CommitmentNormal				337.5** (134.8)	316.5** (135.2)
Constant	570.4*** (86.0)	590.4*** (95.8)	808.8*** (99.3)	413.13*** (62.0)	495.7*** (63.5)
# Observations	105	100	100	94	94
R-Squared	0.055	0.073	0.081	0.063	0.056

Notes: Coefficients from ordinary least squares regression. Robust standard errors in parentheses. Dependent variable *Deviation Signal* corresponds to the absolute difference between the final estimate and the signal. Dependent variable *Accuracy* corresponds to the absolute difference between the final estimate and the true state of 3000.
Level of significance: $*p < 0.1$, $**p < 0.05$, $***p < 0.01$

We begin by looking at treatments Commitment10 and NoCommitment10. First notice that we replicate our main finding from above. On average, the deviation between final estimates and the additional signal is 546.2 points higher in the Commitment10 treatment, compared to the NoCommitment10 treatment, a substantial and statistically significant effect (see *Table 2* (column 2) and *Figure 2* (upper right panel)).

Turning to the accuracy of final estimates, *Table 2* (column 3) reveals a pronounced and significant negative effect of commitment. The average deviation from the correct solution of 3000 is 537.6 points higher in Commitment10, compared to NoCommitment10. Thus, in our context, commitment to a first estimate substantially reduces the quality of information processing, as captured by less accurate final estimates.

Additional evidence from the BeliefAccuracy treatment suggests, that subjects in

Commitment10 not only underreact to the signal in objective terms (as measured by final estimation accuracy) but also relative to subjects' own valuations of the quality of the signal. When being asked to assess the quality of the average estimate of a group relative to their own estimate, on average subjects state a threshold group size of 4, at which they would prefer to be paid by the group average (see *Table 1*). Moreover, every subject prefers the average estimate of a group of 10 to its own estimate. Thus, in a situation where subjects knew the estimation task and had time to think it through, but did not know the signal realizations, every subject in BeliefAccuracy preferred to be paid by the average estimate of a group of 10 subjects. Simple back of the envelope calculations suggest that every subject in Commitment10 should therefore *at least* put equal weight on the signal than on its own estimate.[15] However, 24 out of 50 subjects in Commitment10 put less than equal weight on the signal relative to their own first estimate, suggesting that subjects in the Commitment10 treatment not only objectively, but also according to subjects' own relative valuations respond too little to the additional information.

Finally, we turn to treatments CommitmentNormal and NoCommitmentNormal. First note that also these treatments replicate our main finding. On average, the deviation between final estimates and the additional signal is 337.5 points higher in the CommitmentNormal treatment, compared to the NoCommitmentNormal treatment. This difference in deviations from the additional signal is statistically significant (see *Table 2* (column 4) and *Figure 2* (lower panel)).

Recall that in CommitmentNormal and NoCommitmentNormal, we provided subjects with more objective information, specifically a random draw from a normal distribution

[15]To illustrate this, we make the simplifying assumption that signals are normally distributed, and, based on results from BeliefAccuracy, we conservatively assume that subjects think that the average of a group of 10 subjects is of equal quality than their own estimate. Suppose now that before seeing the bowl with peas, subjects are holding an uninformative prior, that the first estimate m_1 reflects the mean of a distribution $\mathcal{N}(m_1, \sigma)$ and that the additional information m_2 is drawn from $\mathcal{N}(3000, \sigma)$. Since the variance of both distributions is identical, the Bayesian posterior is given by $\frac{1}{2}m_1 + \frac{1}{2}m_2$, i.e., subjects should put equal weight on the first estimate and the additional signal.

with mean equal to 3000 and a standard deviation of 300. Also in a such a scenario, our findings reveal that commitment to a first estimate significantly reduces the quality of information processing. Specifically, the average accuracy of final estimates is 316.5 points lower with commitment *Table 2* (column 5). Thus, also for information drawn from a known distribution with mean equal to the true state, commitment causes a reduction in the quality of information processing, as reflected by less accurate final estimates.

RESULT 3: *Final estimates in Commitment10 and CommitmentNormal are less accurate relative to NoCommitment10 and NoCommitmentNormal, respectively.*

2 Driving Forces of the Commitment Effect

This section investigates the underlying channels of the commitment effect we identified. We interpret our findings in terms of commitment and consistency. Commitment to a first opinion raises resistance to incorporate new information, due to a desire to be consistent with the first opinion. In the next step, we present results from additional treatments that allow us to rule out two potential alternative explanations for our findings. Section 3.2 then presents evidence from a series of treatment conditions that were designed to uncover internal and social motives behind our main finding.

2.1 Alternative Explanations

The first potential alternative explanation is different levels of effort provision. It could be that (for some reason) subjects in the Commitment treatment provided higher effort in the estimation task, thus obtaining more precise private information and therefore rationally incorporating the additional piece of information to a smaller extent. Notice, however, that our experimental procedures ensured that both with commitment and without, subjects had an equal amount of time to form an estimate, before receiving the signal. Also, in both treatments subjects knew that they had to provide an estimate and did not know that they would receive the informative signal at a later point. Nonetheless, an assumption of differential effort provision could potentially be justified, given that subjects in the Commitment treatment had to write down a first estimate (and knew this while thinking about the correct solution of the task), and therefore might have taken the estimation task more seriously. In comparison, while thinking about the solution of the task, subjects in the NoCommitment treatment did not know (though probably could guess) that they would have to write down an estimate. In other words, knowing that one would have to write down an estimate might have increased the motivation to solve the task correctly in the Commitment treatment.

The second alternative explanation is confirmatory bias. Confirmatory bias describes a tendency to interpret new information as being in favor of or confirming the existing hypothesis.[16] It seems unlikely that our findings could be explained by confirmatory bias, since in all treatments, subjects knew that their task was to provide an estimate about the number of peas an the bowl. More specifically, with commitment and without, subjects had time and incentives to form an own hypothesis prior to receiving the additional signal. Thus, if a confirmatory bias exists in our context, it should be present in both treatments. However, since subjects in treatment NoCommitment did not write down a first estimate prior to receiving the signal, we cannot actually verify that they indeed formed an own hypothesis at this point, potentially leaving some room for confirmatory bias to play a role.

We conducted two robustness treatments, **CommitmentR** and **NoCommitmentR** to further address and rule out the differential effort provision explanation and an additional treatment, **NoCommitmentOwnHyp**, that speaks to the confirmatory bias hypothesis. The goal of the CommitmentR and NoCommitmentR treatments was to ensure that while thinking about an estimate, i.e., while putting effort in the task, the treatments are identical in every respect, and only afterwards the treatment manipula-

[16] Rabin and Schrag (1999) model confirmatory bias and show how it can lead to overconfidence. They also provide a review on the psychological literature on confirmatory bias.

tion occurs. Therefore, we fixed the time subjects could see the bowl with peas in both treatments at five minutes,[17] and kept all the information and instructions provided to subjects, as well as procedures etc. identical until after these five minutes.

The treatments were in most parts identical to treatments Commitment and No-Commitment respectively. In both treatments, subjects knew that they had to form an estimate. Specifically, they were informed that they would be shown the bowl with peas for exactly five minutes, and that afterwards we would distribute an answer sheet (and a pen) were they could write down their response. Thus, while seeing the bowl with peas, i.e., while thinking about the solution of the task, subjects in both treatments were in the exact same situation, rendering differential effort provision unlikely. After the five minutes the bowl was removed, and the experimenter waited for an additional 10 seconds. In the CommitmentR treatment, the experimenter then distributed the pens and answer sheets. Notice that subjects typically provided their estimate within only a few seconds after they had received a pen and an answer sheet, suggesting that they had indeed formed an estimate within the five minutes the bowl was displayed, i.e., while the two robustness treatments were still completely identical from the subjects' perspective. After they provided their response, subjects were given the additional information and, just like in the original treatment, were informed that they could revise their estimate. In the NoCommitmentR condition, we simultaneously distributed the additional information, pens and the answer sheets.

To provide additional evidence against the confirmatory bias hypothesis, we ran an additional condition without commitment, treatment **NoCommitmentOwnHyp**. This control treatment was identical to NoCommitment, except that we asked subjects to raise their hand once they had come up with an estimate. After that, subjects received the informative signal and had time to revise their estimate. Thus, in NoCommitmentOwn-Hyp, all subjects indicated that they have formed an own hypothesis prior to receiving

[17]Notice that we chose the time limit of five minutes based on behaviour from our other treatment conditions, where virtually all subjects had provided an estimate within five minutes.

the informative signal. However, subjects did not write down or state their first estimate, i.e., no commitment was made.

Table 3: *Linear Regressions*

	(1)	(2)	(3)	(4)
	CommitmentR and NoCommitmentR	(No)Commitment and (No)CommitmentR		NoCommitmentOwnHyp and Commitment
	Dependent Variable: Deviation Signal			
1 if Commitment	313.33** (139.0)	328.25*** (97.6)	341.68** (138.4)	324.01** (133.4)
1 if Robustness			-21.68 (111.7)	
1 if Commitment and Robustness			-28.36 (196.14)	
Constant	548.71*** (71.3)	559.77*** (55.79)	570.39*** (86.0)	588.06*** (77.7)
# Observations	99	204	204	101
R-Squared	0.049	0.051	0.053	0.054

Notes: Coefficients from ordinary least squares regression. Robust standard errors in parentheses. Dependent variable *Deviation Signal* corresponds to the absolute difference between the final estimate and the signal. Columns (2) and (3) use observations from treatments Commitment, NoCommitment, CommitmentR and NoCommitmentR.
Level of significance: $*p < 0.1$, $**p < 0.05$, $***p < 0.01$

Results: 144 subjects participated in the three additional treatments. *Table 1* provides summary statistics. We begin by analyzing the (No)CommitmentR treatments. We again replicate our main finding from above. On average, the deviation between final estimates and the additional signal is 313.3 points higher in the CommitmentR treatment, compared to the NoCommitmentR treatment, a statistically significant effect (see *Table 3* (column 1)). The effect size is remarkably similar to that in the original treatments. To more cleanly compare effect sizes, columns (2) and (3) of *Table 3* show the results of a difference-in-difference estimation, where absolute deviation is regressed on a commitment dummy, a robustness dummy, and an interaction term equal to one if subjects were in the robustness treatment with commitment (CommitmentR). The coefficient of the interaction term is insignificant, i.e., we cannot reject the nullhypothesis of zero difference between the effect sizes (see *Table 3* (column 3)). In sum, even though differential effort provision is made unlikely by design, we find a significant effect of commitment on information processing that is similar in size to our original treatments.

Turning to treatment NoCommitmentOwnHyp, we find that our main result continues to hold. On average, the deviation between the final estimate and the signal in the Commitment treatment is 324.0 points higher, compared to NoCommitmentOwnHyp, (see *Table 3* (column 4)). Again, the effect size is very similar to that in the two original treatments, as illustrated by the fact that deviations from the signal in NoCommitmentOwnHyp do not significantly differ from NoCommitment ($p = 0.88$, using linear regression). Thus, even though we ensure that subjects formed a prior hypothesis, we find a pronounced and significant treatment effect, suggesting that confirmatory bias is an unlikely driver of our findings.

RESULT 4: *Confirmatory bias as well as different levels of effort provision are unlikely to drive the effect of commitment on information processing.*

2.2 Commitment: Private or Social Roots?

In a final step we investigate if the impact of commitment on information processing has internal or social roots. In our Commitment treatment, commitment was achieved by writing down the current belief and stating it towards the experimenter. Thus, in as much as the experimenter potentially creates social image effects, both internal and social motives could explain the reluctance to incorporate new information that is created by commitment.

We conducted three additional treatments where we systematically manipulate the social image component of commitment. In treatment **CommitmentNoSI**, we removed any social aspect of commitment. The treatment was identical to the Commitment condition, the only difference being that subjects did not state their first estimate towards the experimenter. In CommitmentNoSI, subjects wrote down a first estimate in private. Once all subjects had written down a first estimate, each subject received the sheet containing the additional information. Thus, commitment was achieved by writing down an estimate, but not by communicating it towards the experimenter.

In treatments **CommitmentSI** and **CommitmentSImoney** we strengthened the social element of commitment. In CommitmentSImoney, we introduced bystanders. A bystander observed the estimation behaviour of always two randomly assigned subjects, and was incentivised to select the subject he or she thought gave the most accurate estimate. Bystanders did not see the bowl with peas, but for both subjects observed their first and final estimate respectively, as well as the additional piece of information. Subjects in CommitmentSImoney were informed of the role of the bystanders and were told that they would earn 3 euros extra if they were selected by their bystander. Thus, subjects were observed by a third party and were given monetary incentives to care about what that third party thinks of them in terms of estimation ability. CommitmentSI was identical to CommitmentSImoney, only that subjects did not receive any extra money in case they were selected by their bystander. Thus, in CommitmentSI we increased the social dimension by making estimation behaviour visible to a third party, but did not induce any monetary reasons for subjects to care about the third party.

Taken together, the three treatments allow us to examine if the commitment effect prevails absent any social component, and if the role of commitment can be strengthened by adding social aspects.

Results: 152 subjects participated in the three treatments (*Table 1* provides summary statistics).

We begin by investigating the strength of the commitment effect absent any social component. Column (1) of *Table 4* reveals that the average absolute deviation between final estimates and the additional piece of information is 394.2 points higher in the CommitmentNoSI treatment, compared to NoCommitment, indicating that purely private motives are sufficient to generate a significant effect of commitment. Turning to exogenously induced increases in social motives, column (2) of *Table 4* regresses the absolute deviation from the additional signal on a set of treatment dummies (capturing Commit-

Table 4: *Linear Regressions*

	(1) CommitmentNoSI and NoCommitment	(2) CommitmentSI(money) and Commitment	(3) CommitmentSI(money) and (No)Commitment and CommitmentNoSI
		Dependent Variable: Deviation Signal	
1 if NoCommitment	-394.16***		-341.68**
	(128.7)		(138.4)
1 if CommitmentSI		-5.27	-5.27
		(146.8)	(146.8)
1 if CommitmentSImoney		100.89	100.89
		(160.7)	(160.7)
1 if CommitmentNoSI			52.48
			(144.67)
Constant (CommitmentNoSI)	964.55***		
	(95.7)		
Constant (Commitment)		912.07***	912.07***
		(108.4)	(108.4)
# Observations	105	152	257
R-Squared	0.083	0.003	0.044

Notes: Coefficients from ordinary least squares regression. Robust standard errors in parentheses. Dependent variable *Deviation Signal* corresponds to the absolute difference between the final estimate and the signal. Column (1) is based on observations from treatments CommitmentNoSI and NoCommitment. Column (2) uses observations from treatments CommitmentSI, CommitmentSImoney and Commitment. Column (3) uses observations from treatments CommitmentSI, CommitmentSImoney, CommitmentNoSI, Commitment and NoCommitment. *Level of significance*: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

mentSImoney and CommitmentSI) and a constant (treatment Commitment). No significant differences between treatments CommitmentSImoney and CommitmentSI and Commitment treatment can be detected.[18] Column (3) of *Table 4* uses observations from treatments CommitmentSI(money) and (No)Commitment and CommitmentNoSI, and again regresses the absolute deviation from the additional signal on a set of treatment dummies and a constant (treatment Commitment). Only deviations in the NoCommitment treatment significantly differ from that in the Commitment treatment. An F-test, does not reject the nullyhypothesis that the coefficients for treatments CommitmentSI, CommitmentSImoney and CommitmentNoSI are jointly zero ($F(3, 252) = 0.2$; $p = 0.89$).

In sum, this set of results suggests that private motives are sufficient to generate

[18]Notice that both treatments generate a significant treatment effect relative to NoCommitment ($p < 0.05$, respectively).

a commitment effect on information processing, while exogenously strengthening social aspects does not further add to the commitment effect.

RESULT 5: *Even absent any social elements of commitment, we identify a significant impact of commitment on information processing. Exogenous increases in social motives do not add to the commitment effect.*

3 Concluding Remarks

We have presented experimental evidence highlighting a new motive why individuals might hold on to (wrong) beliefs, even when being confronted with new and challenging information. We find that after committing to a first belief, individuals are reluctant to incorporate new information. This reluctance causes lower belief accuracy and also generates higher belief divergence among subjects. Our findings are in line with the notion that people want to behave consistently. This desire to be consistent with the belief subjects committed to seems to be driven by internal rather than social motives.

Our results highlight a general pattern of consequences from early actions or statements for information processing. These commitments cause stickiness in behaviour and reduce flexibility to adjust to new information. Given previous commitments, changes of mind are particularly hard to execute as inconsistencies become evident and difficult to conceal. These findings might for example also be of interest for the design of institutions such as committee or jury procedures. Institutionally requested commitments on a certain opinion or intention at an intermediate stage can decrease the quality of final choices as these potentially do not reflect the full level of available information.

Behavior and Inequality Research Institute and University of Bonn

University of Zurich

REFERENCES

Abeler, J., Becker, A. and Falk, A. (2014). 'Representative Evidence on Lying Costs', *Journal of Public Economics*, vol. 113, pp. 96-104.

Bénabou, R. and Tirole, J. (2006). 'Belief in a Just World and Redistributive Politics', *Quarterly Journal of Economics*, vol. 121(2). pp. 699-746.

Benoît, J. and Dubra, J. (2011). 'Apparent Overconfidence', *Econometrica*, vol. 79(5) pp. 1591-1625.

Benoît, J., Dubra, J. and Moore, D. (2015). 'Does the Better-Than-Average Effect Show That People Are Overconfident?: Two Experiments', *Journal of the European Economic Association*, vol. 13(2), pp. 293-329.

Brunnermeier, M. and Parker, J. (2005). 'Optimal Expectations', *American Economic Review*, vol. 95(4), pp. 1092-1118.

Burks, S., Carpenter, J., Goette, L. and Rustichini, A. (2013). 'Overconfidence and Social Signaling', *Review of Economic Studies*, vol. 80(3), pp. 949-983.

Camerer, C., and Lovallo, D. (1999). 'Overconfidence and Excess Entry: An Experimental Approach', *American Economic Review*, vol. 89(1), pp. 306-318.

Charness, G., Rustichini, A. and van de Ven, J. (2014). 'Self-confidence and Strategic Deterrence', Working Paper, University of Amsterdam.

Cialdini, R.. (1984). *Influence, the Psychology of Persuasion*, New York: Harper Collins.

DellaVigna, S. and Malmendier, U. (2006). 'Paying Not to Go to the Gym', *American Economic Review*, vol. 96(3), pp. 694-719.

Deutsch, M. and Gerard, H. (1955). 'A study of normative and informational social influences upon individual judgment', *The Journal of Abnormal and Social Psychology*, vol. 51(3), pp. 629-636.

Dohmen, T. and Falk, A. (2011). 'Performance Pay and Multi-dimensional Sorting: Productivity, Preferences and Gender', *American Economic Review*, vol. 101(2), pp. 556-590.

Eil, D. and Rao, J. (2011). 'The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself', *American Economic Journal: Microeconomics*, vol. 3(2), pp. 114-138.

Ewers, M. and Zimmermann, F. (2015). 'Image and Misreporting', *Journal of the European Economic Association*, vol. 13(2), pp. 362-380.

Falk, A., Huffman, D. and Sunde, U. (2006). 'Self-Confidence and Search', Working Paper, University of Bonn.

Falk, A. and Zimmermann, F. (forthcoming). 'Consistency as a Signal of Skills', *Management Science*.

Festinger, L. (1957). *A Theory of Cognitive Dissonance*, Palo Alto CA: Stanford University Press.

Fischbacher, U. and Föllmi-Heusi, F. (2013). 'Lies in Disguise: An Experimental Study on Cheating', *Journal of the European Economic Association*, vol. 11(3), pp. 525-547.

Freedman, J. and Fraser, S. (1966). 'Compliance Without Pressure: The Foot in the Door Technique', *Journal of Personality and Social Psychology*, vol. 4(2), pp. 195-202.

Gneezy, U. (2005). 'Deception: The role of consequences', *American Economic Review*, vol 95(1), pp. 384-394.

Gneezy, U., Rockenbach, B. and Serra-Garcia, M. (2013). 'Measuring lying aversion', *Journal of Economic Behavior and Organization*, vol. 93, pp. 293-300.

Greiner, B. (2003). 'An Online Recruitment System for Economic Experiments', in (K. Kremer and V. Macho, eds.), *Forschung und Wissenschaftliches Rechnen. Gesellschaft für wissenschaftliche Datenverarbeitung Bericht*, pp. 79-93, Göttingen, Germany: Gesellschaft für wissenschaftliche Datenverarbeitung.

Heider, F. (1946). 'Attitudes and Cognitive Organization', *Journal of Psychology*, vol. 21, pp. 107-112.

Kahan, D. (2012). 'Why we are poles apart on climate change', *Judgment and Decision Making*, vol. 8, pp. 407-424.

Kahan, D. (2013). 'Ideology, Motivated Reasoning, and Cognitive Reflection: An Experimental Study', *Nature*, vol. 488, pp. 255.

Kőszegi, B. (2006). 'Ego Utility, Overconfidence, and Task Choice', *Journal of the European Economic Association*, vol. 4(4), pp. 673-707.

Malmendier, U. and Tate, G. (2008). 'Who Makes Acquisitions? CEO Overconfidence and the Market's Reaction', *Journal of Financial Economics*, vol. 89(1), pp. 20-43.

Moebius, M., Niederle, M., Niehaus, P. and Rosenblat, T. (2013). 'Managing Self-Confidence: Theory and Experimental Evidence', Working Paper, Stanford University.

Newcomb, T. (1953). 'An Approach to the Study of Communicative Acts', *Psychological Review*, vol. 60, pp. 393-404.

Piketty, T. (1995). 'Social Mobility and Redistributive Politics', *Quarterly Journal of Economics*, vol. 110(3), pp. 551-584.

Rabin, M. and Schrag, J. (1999). 'First Impressions Matter: A Model of Confirmatory Bias', *Quarterly Journal of Economics*, vol. 114(1), pp. 37-82.

Sherman, S. (1980). 'On the Self-Erasing Nature of Errors of Prediction', *Journal of Personality and Social Psychology*, vol. 39(2), pp. 211-221.

Svenson, O. (1981). 'Are We All Less Risky and More Skillful Than Our Fellow Drivers?', *Acta Psychologica*, vol. 47, pp. 143-148.