

IZA DP No. 7372

Correlation Neglect in Belief Formation

Benjamin Enke
Florian Zimmermann

April 2013

Correlation Neglect in Belief Formation

Benjamin Enke

BGSE, University of Bonn

Florian Zimmermann

*BGSE, University of Bonn
and IZA*

Discussion Paper No. 7372

April 2013

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Correlation Neglect in Belief Formation^{*}

Many information structures generate correlated rather than mutually independent signals, the news media being a prime example. This paper shows experimentally that in such contexts many people neglect these correlations in the updating process and treat correlated information as independent. In consequence, people's beliefs are excessively sensitive to well-connected information sources, implying a pattern of "overshooting" beliefs. Additionally, in an experimental asset market, correlation neglect not only drives overoptimism and overpessimism at the individual level, but also affects aggregate outcomes in a systematic manner. In particular, the excessive confidence swings caused by correlated signals give rise to predictable price bubbles and crashes. These findings are reminiscent of popular narratives according to which aggregate booms and busts might be driven by the spread of "stories". Our results also lend direct support to recent models of boundedly rational social learning.

JEL Classification: C91, D03, D83, D84, D40

Keywords: beliefs, correlation neglect, experiments, markets, overshooting

Corresponding author:

Florian Zimmermann
University of Bonn
Bonn Graduate School of Economics
Kaiserstr. 1
53113 Bonn
Germany
E-mail: florian.zimmermann@uni-bonn.de

^{*} Financial support from the Russell Sage Foundation (Small Grants Program in Behavioral Economics), the International Foundation for Research in Experimental Economics (Small Grants Program) and the Bonn Graduate School of Economics is gratefully acknowledged. For valuable discussions we are grateful to Steffen Altmann, Anke Becker, Thomas Dohmen, Hans-Martin von Gaudecker, Armin Falk, Paul Heidhues, Rosemarie Nagel, Matthew Rabin, Joel Sobel, Hal Varian, Georg Weizsäcker, and Matthias Wibral. Helpful comments were also received from seminar audiences at Amsterdam, WZB Berlin, Bonn, Frankfurt, Maastricht, Pompeu Fabra, and at the 2012 North-American ESA conference in Tucson.

1 Introduction

A pervasive feature of information structures is that decision-makers are exposed to correlated signals.¹ For example, different news media frequently share a common information source such as press agencies, so that the contents of different news reports (newspaper articles, television shows, online print) tend to be correlated. Similarly, in social networks, opinions from different network members are often partly based on information from a mutually shared third party, implying a correlation between people’s views. Hence, in communicating with these people, one is confronted with correlated information. A common feature of these examples is that, as a result of the information structure, similar “stories” are getting told and retold multiple times. However, little is known about the process of belief formation in the presence of correlated signals. Do people fully understand the consequences of such information structures? If not, what are systematic economic implications of this, both at an individual and possibly at an aggregate level?

Taking these questions as starting point, the contribution of this paper is twofold. First, we use a simple experiment to provide clean evidence that even in transparent settings many people neglect correlations in information sources when forming beliefs. In consequence, people’s beliefs are excessively sensitive to influential “base signals” (well-connected information sources), so that posteriors follow a systematic “overshooting” pattern. In a second step, we examine the aggregate implications of such correlation neglect. Recently, Shiller (2000) and Akerlof and Shiller (2009) have argued that “exuberant” public opinions or “panics”, driven by the “spread of stories”, might be a driver of aggregate distortions. In this spirit, we investigate whether the incidence of correlated (and hence similar) news stories affects aggregate outcomes in a predictable manner. We demonstrate that, in an experimental asset market, correlation neglect leads to pronounced and predictable price bubbles and crashes.²

In the baseline experiment, subjects need to estimate an ex ante unknown state of the world and are paid for accuracy. Subjects are provided with computer-generated, unbiased (but noisy) information about the true state of the world. The key idea of our experimental design is to construct two sets of information (one with and one without a known and simple correlation), which result in the same Bayesian posterior. In a between-subjects design, one group of subjects receives correlated,

¹In the following, a correlation is implicitly understood as being conditional on a state realization. Also, we only refer to positive correlations.

²For the purpose of this paper, we define a bubble (crash) to be a situation in which prices are significantly above (below) fundamental values.

the other uncorrelated information. If subjects correctly take the correlations in the correlated treatment into account, estimates should be identical across treatments. However, our results indicate that a considerable fraction of subjects treats all incoming information as approximately independent and hence exhibits correlation neglect. Consequently, beliefs in the correlated information structure overshoot in the predicted direction, an effect that is sizeable, significant, and causes lower payoffs. Thus, while beliefs remain statistically unbiased *ex ante*, they are highly sensitive to certain types of signals and exhibit excessive swings: If a relatively low (high) base signal realizes, the repeated emergence of this signal through other signals leads many people to “double-count” the base signal and hence to become overpessimistic (overoptimistic). We further document considerable heterogeneity in people’s propensity to disregard correlations and therefore investigate the sources of the bias. While a substantial increase in the stake size did not induce subjects to develop more accurate beliefs, the bias is significantly associated with cognitive ability. These findings suggest that neglecting correlations is driven by subjects’ limited understanding of the implications of correlations as opposed to a “rational” tradeoff between higher earnings and lower cognitive effort costs.

An important question is whether these biased beliefs have systematic implications beyond the individual level. To approach this issue, we embed our individual belief elicitation design into a continuous double-auction environment in which subjects trade financial assets. The fundamental value of the assets equals the true state from the baseline design. Before each trading round, subjects receive the same sets of information about the true state as in the individual treatments. Again, we form treatment (control) groups by providing correlated (uncorrelated) signals about the fundamental value of the assets. The results show that market prices differ between treatments in the direction one would expect if subjects disregard correlations. In rounds in which correlation neglect leads to overly optimistic beliefs, market prices are too high relative to both the control treatment and the fundamental level. Likewise, in rounds in which neglecting correlations implies pessimism, market prices are too low. Thus, correlation neglect predictably drives price bubbles and crashes, respectively, implying excessively volatile price levels. In addition, a trader’s propensity to ignore correlations predicts his asset holdings, leading to excessive buying during bubbles and very little buying during crashes. Consequently, traders with a higher tendency to overlook correlations earn significantly lower profits. By emphasizing the influence of correlation neglect on aggregate outcomes, our findings also demonstrate its robustness to market interaction (for other studies investigating the implications of individual biases in markets

see e.g. Camerer (1987) and Gneezy, Kapteyn and Potters (2003)).

In order to gain clean evidence on the cognitive bias under study, we deliberately chose a tightly controlled, abstract decision environment. In the last part of the paper, we demonstrate evidence for correlation neglect in a more natural environment in which subjects extract information from newspaper articles. Subjects are asked to estimate the growth of the German economy in 2012 and are paid for accuracy. We avail ourselves of a particular correlation between two forecasts of different research institutes and provide subjects with two newspaper articles reporting these forecasts. In the control condition, subjects receive the same information in an uncorrelated way. Despite the salient correlation between the forecasts, we again find significant evidence for correlation neglect. This highlights the robustness of our findings and suggests that people hold distorted beliefs even in familiar informational settings.

The results in this paper provide important insights into the nature of belief formation and corresponding implications for aggregate variables. In particular, correlation neglect predictably produces overshooting beliefs, i.e., overoptimism and overpessimism, which carry over to market outcomes such as bubbles, crashes, and volatile price levels. Notably, these findings provide a micro-foundation for long-standing claims that from time to time markets undergo periods of “irrational exuberance” or “excessive doubt”. This suggests that excessive movements in public confidence might be driven by the ubiquity of correlated information which implies that people are repeatedly confronted with similar stories. Such an interpretation of our results is consistent with the arguments provided by Akerlof and Shiller (2009) who emphasize how the telling and retelling of stories drives confidence swings. While different psychological theories can be invoked to explain either (collective) overoptimism or -pessimism, correlation neglect provides a unified view on these phenomena and relates them to the informational network structure.³

The results from our individual belief elicitation treatments also admit a very natural interpretation in terms of social learning in networks. Our static experimental environment with exogenous signals allows for a clean identification of people’s tendency to ignore correlations when updating beliefs and hence provides direct empirical support for recent models of “naïve” social learning. Based on the work by DeGroot (1974), the contributions by, inter alia, DeMarzo, Vayanos and Zwiebel (2003) and Golub and Jackson (2010) analyze social network models of opinion formation in which the agents suffer from “persuasion bias”. In these models, agents

³Wishful thinking (overoptimism) theories of distorted beliefs in markets are provided by Brunnermeier and Parker (2005) and Bénabou (2013).

heuristically treat all reports of their network neighbors as independent, even though these opinions might be partly based on identical (third party) information. The authors show that this assumption leads to the phenomenon of social influence, meaning that an agent’s impact on group opinions depends on how connected he is within the network. The updating rule of thumb which generates these results can naturally be understood as an extreme form of correlation neglect. Similarly, in the sequential herding literature, Eyster and Rabin (2010) develop a model of naïve herding in which agents believe that each predecessor’s action only reveals that agent’s private information. In essence, by disregarding the correlation among their pre-movers’ actions, agents thus double-count early signals and fall prey to correlation neglect. This form of naïveté predicts herding behavior in a broad set of environments where rational agents would never herd. Thus, in both types of social learning models, (collective) beliefs are excessively sensitive to some base signals. One contribution of our paper is to provide direct evidence on the validity of the underlying assumption of this type of inferential naïveté.⁴

A key insight of this paper is that, in the presence of correlated signals, people’s beliefs tend to overshoot. Since correlated information often arises through mere communication in multi-person settings, our basic finding could also shed light on a well-known puzzle in the group decision-making literature. According to the so-called group polarization effect (first established in experiments summarized by Stoner (1968)), groups tend to hold more extreme views than a summary statistic of the beliefs of all group members. If each group member can be understood to have received some signal about the matter in question (Sobel, 2012) and if these signals are correlated across people (as is the case in corresponding experiments where all group members often receive the same information), then correlation neglect predicts group polarization. In essence, if each group member treats the opinions of others as independent signals, the group will double-count the available information and hence overshoot in its collective decision.

Finally, our paper relates to findings in Kallir and Sonsino (2010) and Eyster

⁴There is some indirect evidence in the experimental herding literature pointing to naïve social learning. First, Penczynski (2012) finds that many experimental players can be classified as level $k = 2$ thinkers, which in this context means that people best respond to the belief that all other players’ actions are based on their respective private signal only. Second, in a meta study, Kübler and Weizsäcker (2005) argue that many people fail to recognize herding behavior of others and thus attribute too much informational content to each pre-mover’s action. This insight is not driven by lack of common knowledge of rationality, since it also holds if all predecessors are computers with known (rational) decision rules (Grebe, Schmid and Stiehler, 2008). Corazzini, Pavesi, Petrovich and Stanca (2012) explicitly test the persuasion bias hypothesis in a social network setup, but find only mixed results. Since their focus is on consensus (“limit”) beliefs, their design does not allow for a clean identification of correlation neglect. Loosely related is also the literature on persuasive communication, see DellaVigna and Gentzkow (2010) for a survey.

and Weizsäcker (2011). These contributions show experimentally that people have difficulties accounting for correlations in asset returns when constructing financial portfolios. For example, Eyster and Weizsäcker (2011) devise pairs of portfolio allocation problems in which the correlation among the asset returns differs within pairs, yet rational decision makers would always choose the same portfolio. They find that allocations differ within pairs in a way which suggests that people neglect the correlation structure among the returns. While the cognitive mechanism underlying these findings might be similar to the one we investigate, correlation neglect in information sources and the resulting belief structure constitute conceptually different phenomena.

The remainder of the paper is organized as follows. In the next section, we introduce a simple formal framework of correlation neglect. In section 3, we present our main experiment including the market treatments. Section 4 contains our newspaper study and section 5 concludes.

2 Conceptual Framework

In this section, we present an analytical framework which fleshes out the implications of correlation neglect for posterior beliefs. We use a simple set-up with uninformative priors and two signals to formalize these effects and to introduce the terminology to be employed throughout the paper. The insights easily generalize to more generic settings, in particular to settings where decision makers do have prior beliefs.

Suppose an unknown state of the world is characterized by the realization of a random variable $\theta \sim \mathcal{N}(\mu, \sigma_\theta^2)$. Nature generates unbiased and independent signals s_i , $i \in \{1, 2\}$, with $s_i = \theta + \epsilon$, $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$, so that the signals are of identical quality. A decision maker needs to estimate θ . However, instead of observing s_1 and s_2 , he observes s_1 and

$$\tilde{s}_2 = \alpha s_1 + (1 - \alpha)s_2, \quad \alpha \in (0, 1),$$

i.e., the two signals are correlated conditional on θ . We will also refer to the first signal (s_1) as the base signal. Suppose that α is a known parameter. Further assume for simplicity that the decision maker's prior distribution about θ is uninformative.⁵ After observing the two signals, a rational agent fully extracts s_2 from \tilde{s}_2 and computes the average of the two independent signals. His beliefs are then given

⁵Technically, the prior is normally distributed with variance τ^2 , where $\tau^2 \rightarrow \infty$.

by the (Bayesian) posterior distribution $\mathcal{B} = \mathcal{N}(m_B, \sigma_B^2) = \mathcal{N}(\frac{s_1+s_2}{2}, \frac{\sigma_\epsilon^2}{2})$. Now suppose that the decision maker suffers from correlation neglect, i.e., he does not fully take into account the extent to which \tilde{s}_2 reflects s_1 , but rather takes \tilde{s}_2 (to some extent) at face value. Call such a decision maker naïve and let his degree of naïveté be parameterized by $\chi \in [0, 1]$ such that $\chi = 1$ implies full correlation neglect. Throughout the paper, we will refer to agents with $\chi = 0$ as Bayesian or sophisticated agents. A naïve agent extracts the second independent signal from \tilde{s}_2 according to the rule

$$\hat{s}_2 = \chi\tilde{s}_2 + (1 - \chi)s_2 = s_2 + \chi\alpha(s_1 - s_2)$$

where \hat{s}_2 denotes the agent's (possibly biased) inference of s_2 . That is, like a Bayesian agent, a naïve agent attempts to extract the underlying independent signal, yet only extracts a linear combination of s_2 and the observed signal \tilde{s}_2 due to his naïveté. He thus forms posterior beliefs according to the correlation neglect distribution $\mathcal{CN} = \mathcal{N}(m_{CN}, \sigma_{CN}^2) = \mathcal{N}(\frac{s_1+\hat{s}_2}{2}, \frac{\sigma_\epsilon^2}{2})$.⁶

Prediction. *Assuming that $\chi > 0$, following a sequence of correlated signals, beliefs exhibit an overshooting pattern. Given a high (low) base signal, i.e., $s_1 > s_2$ ($s_1 < s_2$), the beliefs of a naïve agent are biased upward (downward) compared to the Bayesian benchmark.*

Intuitively, by partially neglecting the correlation between the two signals, the decision-maker double-counts the first signal, so that the posterior belief is systematically biased in the direction of the first signal. Throughout the paper, we will call a belief above (below) the Bayesian belief overoptimistic (overpessimistic).

Note, however, that the beliefs of a naïve agent remain statistically unbiased. Since the base signal is unbiased, any double-counting leads to a zero expected error. Nevertheless, conditional on any realization $s_1 \neq \theta$, beliefs will overshoot. The upshot of this is that naïve agents are correct on average, yet exhibit excessive swings in their beliefs.

Appendix A1 shows that this bias does not vanish if the number of signals grows large. To the contrary, conditional on some $s_1 \neq \theta$, the expected bias of a naïve agent increases in the number of (correlated) signals, so that even very information-rich environments produce biased beliefs.

⁶While a naïve agent assesses the variance of his belief to be $\sigma_{CN}^2 = \frac{\sigma_\epsilon^2}{2}$, the true variance of his posterior distribution equals $\sigma_\epsilon^2 \left(\frac{1+\chi^2\alpha^2}{2} \right) > \sigma_{CN}^2 \Leftrightarrow \alpha\chi > 0$. That is, in a correlated information structure, an agent will be overconfident about the precision of his belief if and only if he is naïve.

3 Evidence for Correlation Neglect in Belief Formation

We developed a simple experimental design which allows for both the clean identification of correlation neglect and the investigation of its implications in market settings in a unified and coherent framework. This integrated approach will allow us to cleanly identify how the individual bias translates into market outcomes. We first describe the basic belief elicitation design and then explain how these individual treatments were extended into market treatments. After stating our predictions, we present the results.

3.1 Experimental Design

3.1.1 Individual Treatments

An environment in which updating from correlated sources can be studied requires (i) control over signal precision and correlation, (ii) a control condition that serves as a benchmark for updating in the absence of correlated information, and (iii) incentivized belief elicitation. Moreover, the setup should be as simple as possible to allow subjects to specifically focus on the information structure.

Our design accommodates all these features. Subjects were asked to estimate an ex ante unknown continuous state of the world μ and were paid for accuracy. The task was framed as guessing how many items are contained in an imaginary container. In order to keep the experiment as simple as possible, we refrained from inducing prior beliefs. Subjects neither saw the container nor the items it contained. The only information provided to participants consisted of unbiased (but noisy) computer-generated signals about the true state. The key idea of the between-subjects design was to construct two sets of signals (one with and one without a known and simple correlation), which result in the same Bayesian posterior. As depicted in Figure 1, subjects in the correlated treatment received correlated information, while subjects in the control condition received uncorrelated information about μ .

The four computers generate four iid signals about μ . Technically, this was implemented by random draws from a truncated discretized normal distribution with mean μ and standard deviation $\sigma = \mu/2$.⁷ In the correlated treatment (left panel), the intermediaries 1 to 3 observe both the signal of computer A and of computers B to D, respectively. Then they report the average of these two signals.

⁷Truncation was at $\mu \pm 2\sigma = \mu \pm \mu$ in order to avoid negative signals.




Figure 1: Correlated (left panel) and uncorrelated (right panel) information structure

Subjects were provided with information from computer A as well as from the three intermediaries. For example, in one estimation task, the estimates of computers A through D were given by 12, 9, 10, and 0, respectively. Accordingly, each of the three intermediaries took the average of 12 and the corresponding signal of the other computer it communicated with. Thus, computer A reported 12, intermediary 1 reported 10.5, intermediary 2 reported 11, and intermediary 3 reported 6.

Since subjects knew computer A’s signal, they could extract the other independent signals from the intermediaries’ reports. Note that this structure mirrors the examples provided in the introduction. For example, one could think of computer A as a press agency which sells information to various newspapers, which in turn each have an additional independent information source. Alternatively, in a social learning context, the intermediaries could be viewed as network members who each received an independent piece of information, yet have all also talked to a common friend (computer A) before communicating their opinion. In the uncorrelated treatment (right panel), the intermediaries 1 to 3 directly transmitted the signals of computers B to D, respectively. Again, subjects received information from computer A as well as from the three intermediaries. Continuing the example from above, computer A reported 12, intermediary 1 reported 9, intermediary 2 reported 10, and intermediary 3 reported 0.

Note that our identification strategy relies solely on the identical informational content of the two sets of signals. Differences in estimates between the correlated and uncorrelated condition can only be attributed to variations in the information structure since all other factors are held constant. In particular, our design makes no assumptions on prior beliefs.

Upon receiving the four pieces of information, a subject had five minutes to provide his estimate. Subjects completed a total of ten independent estimation tasks without feedback between rounds. We used three different randomized orders of rounds. At the end of the experiment, one round was randomly chosen and

Table 1: Overview of the estimation rounds

True state	Computer A	Intermed. 1 uncorr.	Intermed. 2 uncorr.	Intermed. 3 uncorr.	Intermed. 1 corr.	Intermed. 2 corr.	Intermed. 3 corr.	Bayesian belief	Correlation neglect belief
10	12	9	10	0	10.5	11	6	7.75	9.88
88	122	90	68	5	106	95	64	71.25	96.63
250	179	295	288	277	237	234	228	259.75	219.38
732	565	847	650	1'351	706	608	958	853.25	709.13
1'000	1'100	1'060	629	1'100	1'085	870	1'105	974.75	1'042.38
4'698	1'608	7'240	4'866	5'526	4'424	3'237	3'567	4'810.00	3'209.00
7'338	9'950	1'203	11'322	11'943	5'577	10'636	10'947	8'604.50	9'277.25
10'000	2'543	10'780	6'898	8'708	6'662	4'721	5'626	7'232.25	4'887.63
23'112	15'160	21'806	20'607	47'751	18'483	17'884	31'456	26'331.00	20'745.50
46'422	12'340	32'168	49'841	61'293	22'254	31'091	36'817	38'910.50	25'625.25

The reports of intermediaries 1 through 3 in the uncorrelated condition directly reflect the draws of computers B-D. The Bayesian belief is computed by taking the average of the estimates of computers A-D. The correlation neglect belief is computed assuming $\chi = 1$, i.e., full correlation neglect. Thus, this benchmark is given by the average of the estimate of computer A and the reports of intermediaries 2-4 in the correlated condition. Note that subjects faced the ten rounds in randomized order.

subjects were paid according to the precision of their estimate in that round using a quadratic scoring rule (Selten, 1998). Table 1 provides an overview over the ten rounds. In order to provide an indication of both the direction and the extent of a potential bias, we also provide the benchmarks of full correlation neglect and Bayesian beliefs, respectively.⁸

Subjects received extensive written instructions which explained the details of the task and the incentive structure. In particular, the estimates of the four computers, how these estimates map into the reports of the intermediaries, and the fact that the four computers are of identical quality, were explained in great detail. The instructions included the applicable panel from Figure 1 and it was recommended to subjects to place this figure in front of them when determining their estimates. The instructions also contained an example consisting of estimates of the four computers as well as the respective reports of the three intermediaries, given a certain state of the world. A summary of the instructions was read out aloud. In addition, subjects completed a set of control questions with a particular focus on the infor-

⁸The Bayesian belief is defined as average of the estimates of computers A to D. The correlation neglect belief is defined as average of the estimates of computer A and of the intermediaries 1 to 3 (in the correlated treatment). Thus, the computation of these posteriors ignores both subjects' potential priors and the truncation in the signal distribution. Note that the quantitative errors resulting from this are likely to be very small in magnitude. First, since we do not provide any information on the specific estimation tasks, prior beliefs (if present) are very likely to be vague. Second, the error resulting from ignoring the truncation is small since the tails outside the truncation are fairly thin. Moreover, our definition of the posterior beliefs conforms with observed behavior in the uncorrelated treatment, where subjects tended to merely take the average of the four signals. Finally, and most importantly, these definitions of the posterior beliefs have no effect on the qualitative predictions of our treatment comparison.

mation structure. For example, in both treatments, subjects had to compute the reports of intermediaries 1 and 2 given exemplary estimates of the four computers in order to make sure that subjects understood the (un)correlated nature of the signals.⁹

At the end of the experiment, we conducted a questionnaire in which we collected information on sociodemographics such as gender, age, and monthly disposable income. We also elicited a measure of risk aversion. To capture dimensions of cognitive ability, we asked subjects for their final high school grade (German “Abitur”) and had them solve ten rather difficult IQ test Raven matrices.¹⁰

Finally, we investigated whether neglecting correlations is sensitive to the financial incentives to hold accurate beliefs. To this end, for both treatments, we administered a high-stakes as well as a low-stakes condition. In the high-stakes condition, maximum variable earnings were 30 euros, while in the low-stakes condition subjects could earn at most 10 euros.¹¹

3.1.2 Market Treatments

In the market treatments, the estimation task was embedded into a double-auction setting with uncertainty over the value of the assets. More specifically, in each trading round, an asset’s value corresponded to the true state of the world from the individual treatments. Before each round, traders received the same sets of signals about the true state as participants in the baseline design (see Table 1). In the correlated market treatment, all market participants received correlated, in the control market treatment they received uncorrelated information. Before each trading round, subjects were given five minutes to think about an asset’s value and to provide a non-incentivized estimate. Afterwards, subjects traded the assets.

In order to keep the experiment as simple as possible and to retain the subjects’ focus on the information structure, participants were assigned to be in the role of a buyer or a seller, so that each subject could either buy or sell assets, but not both. A market group consisted of four buyers and four sellers. Subjects were randomly assigned to be in either role and kept their roles throughout the experiment. Subjects remained in the same market groups throughout the ten trading rounds. Before each of the ten rounds, each seller was endowed with four assets. Also, at

⁹The experimental instructions including the control questions are available upon request.

¹⁰In order to learn whether correlation neglect is associated with other established updating mistakes, we also measured the degree to which subjects are vulnerable to base-rate neglect and the gambler’s fallacy. However, we did not find any significant associations.

¹¹Earnings in euros were given by $\pi = 10 - 160 \times (\text{Estimate} / \text{True state} - 1)^2$ in the low-stakes and by $\pi = 30 - 480 \times (\text{Estimate} / \text{True state} - 1)^2$ in the high-stakes condition.

the beginning of each round, each buyer received a monetary endowment that was sufficient to purchase between three and five assets at fundamental values.¹²

Buyers and sellers traded with each other in a standard double-auction format. Buyers could post buying prices and accept selling offers from the sellers. Sellers could post selling prices and accept buying offers from the buyers. Buying and selling offers were induced to converge by the standard procedure, i.e., a new buying (selling) offer had to be higher (lower) than all previous offers. An accepted offer implied a trade and erased all previous offers. Trading lasted for four minutes. Profits per trading period for both buyers and sellers corresponded to the value of the assets owned plus the amount of money held at the end of the respective trading round minus some known fixed costs.

Subjects completed ten trading rounds. We used two different randomized orders of rounds. After each round, subjects received feedback about the true state of the world and the resulting profits from their trading activities in that round. At the end of the experiment, one of the ten rounds was randomly selected and implemented, i.e., payoff-relevant for the subjects. The written instructions included the same information on the information structure as in the individual treatments. A summary of the instructions was read out aloud. In addition to the standard control questions about the information structure, we asked several questions related to the trading activities. After the control questions, there was a test round after which the participants again had the opportunity to ask questions.

3.2 Procedural Details

The experiments were conducted at the BonnEconLab of the University of Bonn. Subjects were mostly students from the University of Bonn and were recruited using the online recruitment system by Greiner (2004). No subject participated in more than one session. The experiment was run using the experimental software z-Tree (Fischbacher, 2007). A total of 140 subjects participated in the individual treatments, in which sessions lasted about 1.5 hours. In the low-stakes condition, average earnings of the 70 subjects equalled 10.80 euros, while subjects in the high-

¹²Throughout the experiment, profits, prices etc. were described in points rather than euros. Since the true state differed in magnitude from round to round, we had to adjust the point / euro exchange rate between rounds. This was made clear to subjects in the instructions. In principle, the exchange rate as well as the budget was informative of the true state. However, the relationship between these variables was chosen to be non-constant across rounds, so that the informational content was rather weak (see the instructions for details). Accordingly, the number of assets a buyer could purchase with his endowment was not constant across rounds. In any case, since budgets and exchange rates were identical across treatments, this procedure cannot explain potential treatment differences.

stakes condition earned 26.30 euros on average. 288 subjects participated in the market treatments. These sessions lasted about 2.5 hours and subjects earned 19.40 euros on average. In all treatments, payments included a 6 euros show-up fee.

3.3 Predictions

Since the informational content of the signals was identical across treatments, subjects should in principle hold the same beliefs in the correlated and the uncorrelated conditions. If, however, subjects neglect correlations, then naïve subjects in the correlated treatment will double-count the base signal of computer A and hence hold systematically biased (overshooting) beliefs relative to the uncorrelated condition. Following a relatively high (low) base signal, such subjects will be overoptimistic (overpessimistic), see also section 2.

In the market treatments, the standard theoretical prediction is that the competitive market equilibrium price is given by the Bayesian belief. Since it is well-established that double-auctions tend to converge to the theoretical competitive equilibrium, this is also the standard experimental prediction. However, if a sufficient number of traders in the correlated treatment exhibits some degree of correlation neglect, the theoretical competitive equilibrium price will be distorted.¹³ In particular, depending on whether the base signal is relatively high or low, one would observe price bubbles or crashes, so that, in essence, aggregate outcomes would reflect the individual belief bias.

On the other hand, it has been argued that the influence of cognitive biases on aggregate variables is limited. In the market environment we implement, two channels in particular might attenuate such effects. First, competitive forces and market incentives could induce subjects to think harder and thus cause a reduction of correlation neglect. Second, markets provide opportunities for traders to learn, both from the actions of more rational market participants and from feedback resulting from profits. While all these channels could mitigate the effect of individual biases on market outcomes, the learning arguments in particular would suggest that correlation neglect (and its consequences) is reduced in the last trading rounds.¹⁴

¹³To see this, heuristically, suppose that a fraction α fully ignores correlations and a fraction $1 - \alpha$ holds Bayesian beliefs. Further suppose that each seller owns four assets and each buyer has a budget sufficient to buy four assets at fundamental values. Then, assuming that subjects do not learn from others' trading behavior and are risk-neutral, the supply and demand curves will be step functions which overlap at the correlation neglect belief if $\alpha \rightarrow 1$. Similar arguments apply if a fraction α exhibits only partial (or heterogeneous degrees of) correlation neglect.

¹⁴Camerer (1987) provides a more extensive discussion of these feedback and learning effects. Similar to our approach, he uses experimental markets to test if other updating mistakes (e.g. base-rate neglect) matter for market outcomes. See also Forsythe, Nelson, Neumann and Wright

3.4 Results

The analysis of our results proceeds in two steps. First, we provide evidence for consistent correlation neglect across the ten estimation tasks in the individual treatments and investigate the sources of this bias. In a second step, we show that correlation neglect persists in markets and leads to volatile price levels.

3.4.1 Individual Treatments - Correlation Neglect and its Sources

First note that varying the stake size had no effect on subjects' beliefs. Thus, in what follows, we pool the results from the high-stakes and the low-stakes condition. At the end of this section, we return to a formal discussion of the consequences of variations in the financial incentives.

Beliefs Across Treatments

Result 1. *Beliefs consistently differ between treatments in the direction predicted by correlation neglect.*

Table 2 provides summary statistics and presents the key findings for all ten estimation rounds. In particular, the table provides median beliefs for both treatments as well as p-values of Wilcoxon ranksum tests.¹⁵

Note that estimates in the uncorrelated condition follow the Bayesian posteriors very closely. Thus, when receiving four independent signals, subjects end up with almost unbiased and rather precise estimates. In contrast, median beliefs in the correlated condition always lie between the Bayesian and the full correlation neglect predictions. Accordingly, Table 2 shows that estimates in all ten estimation tasks differ between treatments in the direction predicted by correlation neglect. In all rounds but one, this difference is statistically significant (Wilcoxon ranksum test). The fact that estimates in the correlated treatment are consistently further away from the Bayesian posterior than estimates in the control condition is also reflected by lower earnings of subjects in the correlated treatment. On average, in the high-stakes (low-stakes) condition, these subjects earned roughly seven (two) euros less than those in the control group, which amounts to 35 (45) % of subjects' average

(1992), Ganguly, Kagel and Moser (2000), and Kluger and Wyatt (2004) for similar studies.

¹⁵In almost all rounds, we observe a few outlier estimates which are hard to rationalize given the information subjects received in the respective rounds. Most of these outliers can be rather clearly attributed to simple typing mistakes, e.g., forgetting a digit or accidentally entering an extra digit. Thus, in order to obtain a clearer view on the cognitive bias, we focus on medians throughout the paper. However, our main results are robust to using means and we report the corresponding test statistics where applicable.

Table 2: Correlation neglect by estimation task

True State	Bayesian Belief	Correlation Neglect Belief	Median Estimate Control Treatment	Median Estimate Correlated Treatment	Ranksum Test (p-value)
10	7.75	9.88	7.95	9.2	0.0002
88	71.25	96.63	71.25	88.15	0.0000
250	259.75	219.38	260	250	0.0064
732	853.15	709.13	850	774	0.0063
1'000	974.75	1'042.38	991	1024	0.0180
4'698	4'810	3'209	4'810	4'500	0.0009
7'338	8'604.5	9'277.25	8'653	9'053.15	0.6983
10'000	7'232.25	4'887.63	7'232	6'395	0.0002
23'112	26'331	20'745.5	25'018	21'555	0.0000
46'422	38'910.5	25'625	38'889	30'054	0.0149

See Table 1 for details of the computation of the Bayesian and the correlation neglect benchmarks. Note that subjects faced the ten rounds in randomized order.

variable earnings. The earnings difference is significant (p-value = 0.0086, Wilcoxon ranksum test).

Table 2 also reveals that correlation neglect distorts beliefs in both directions, i.e., subjects' estimates overshoot upwards or downwards depending on the base signal. Correlation neglect is very stable between rounds and does not seem to depend on the magnitude of the true state. For example, we find evidence for correlation neglect when the true state was 10, but also when the true state was 46'422. Note also that we do not find any order effects, i.e., subjects do not seem to learn to deal with correlations over time. In an OLS regression of all beliefs on a linear and a quadratic time trend, the time regressors are jointly insignificant, see Appendix A2 for details.

In order to be able to study the determinants as well as the implications of correlation neglect at the individual level, we compute a robust measure of an individual's belief "type". This analytical tool is obtained by normalizing the data and aggregating them across rounds into a single belief per individual.

More precisely, for each estimation task and for each individual we first normalize the belief such that it takes on the value zero if it equals the Bayesian posterior and the value one if it equals the full correlation neglect belief. This normalization allows us to compare beliefs across rounds. Moreover, the normalized belief equals the naïveté parameter χ from our conceptual framework.¹⁶ We then compute the median normalized belief over the ten rounds per subject for further analysis.

¹⁶Formally, let $j \in \{1, \dots, 10\}$ denote the estimation task and $i \in \{1, \dots, 140\}$ the individual. The Bayesian posterior in round j is given by b_B^j and the full correlation neglect belief by b_{CN}^j .

Result 2. *Aggregating the data over the 10 estimation rounds confirms that estimates in the correlated treatment are significantly distorted in the direction predicted by correlation neglect.*

Comparing aggregated beliefs across treatments reveals a pronounced and statistically highly significant difference (p-value < 0.0001, Wilcoxon ranksum test). This result is robust to using the mean of normalized beliefs for each individual (p-value = 0.0001, Wilcoxon ranksum test). The treatment difference is visualized by Figure 2, which provides kernel density estimates of the aggregated data, i.e., of the median normalized beliefs.¹⁷ The kernel density reveals that in the uncorrelated treatment the median beliefs of the vast majority of subjects are close to the Bayesian belief, as indicated by the spike around zero. In the correlated treatment, however, we observe two (small) peaks around the Bayesian and the full correlation neglect belief, with many beliefs lying in between.




Figure 2: Kernel density estimates of median normalized beliefs

Sources of Correlation Neglect

Given the substantial belief heterogeneity in the correlated condition evident from Figure 2, we now explore the determinants of falling prey to correlation neglect. We distinguish two explanations for the bias. First, subjects could actually be

The set of normalized beliefs y_i^j is then given by $y_i^j = \frac{b_i^j - b_B^j}{b_{C.N}^j - b_B^j} = \chi_i^j \quad \forall i, j$. Note that this normalization procedure takes into account that the (percentage) difference between Bayesian and correlation neglect belief differs across tasks. Naturally, in the experiment, not all beliefs map into the [0,1] interval. For example, a subject who fully neglects correlations might in addition make a computational mistake to end up with a belief higher than one. Likewise, a subject who aims at computing the Bayesian posterior might make a small error, so that his belief might be below zero.

¹⁷Throughout the paper, we use Silverman’s rule of thumb for bandwidth selection.

incapable of understanding the implications of the information structure and to perform adequate statistical inferences. Second, people could in principle be able to perform the correct computations, yet incur some form of thinking costs and hence “rationally” trade off the benefits of better beliefs against lower cognitive effort costs. According to the latter argument, subjects would actively decide between merely taking the average of the correlated signals on the one hand and performing the more effort-intensive correct calculations on the other hand.

In our context, these explanations have profoundly different implications. If people are in principle capable of performing adequate inferences and choose their cognitive effort level, then an increase in the marginal financial incentives should reduce the amount of correlation neglect. If, on the other hand, subjects try to perform the correct calculations yet miss the cognitive abilities or knowledge to do so, then the stake size should have no effect on the degree of correlation neglect. Rather, the effect should be driven by a limited understanding of statistical correlations. We find suggestive evidence for the latter argument.

Result 3. *While correlation neglect is significantly associated with cognitive ability, the financial incentives to hold accurate beliefs have no impact on subjects’ tendency to disregard correlations.*

Recall that in all previous analyses we pooled the data from our high-stakes and low-stakes conditions. We now provide evidence that tripling the stake size had indeed no effect on subjects’ propensity to ignore correlations. To this end, columns (1) and (2) of Table 3 show the results of a difference-in-difference OLS estimation of each subject’s median normalized belief on (i) a treatment dummy, (ii) a stake size dummy, and (iii) an interaction term equal to one if subjects were in the high-stakes correlated treatment. Recall that a higher normalized belief indicates a stronger propensity to commit correlation neglect. If the increase of the stake size by 200 % leads to more accurate beliefs, then this interaction term should have a negative coefficient. However, the only sizable and significant effect is the treatment difference, which is robust to increasing the (marginal) financial incentives. Interestingly, column (3) shows that the stake size did have a substantial positive impact on subjects’ response times. This suggests that subjects indeed provided higher (cognitive) effort when confronted with higher stakes. However, this higher effort level did not translate into more accurate beliefs. Thus, it seems as if (some) subjects do not know just how to take care of the correlations. For instance, it may be that subjects invest more effort by double-checking their calculations, yet miss the appropriate strategy to develop unbiased beliefs.

Table 3: Correlation neglect and stake size

	Median normalized belief		Median response time
	(1)	(2)	(3)
1 if correlated	0.417*** (0.074)	0.428*** (0.109)	23.55*** (8.770)
1 if high stakes		-0.0469 (0.084)	18.75** (8.770)
1 if correlated high stakes		-0.0207 (0.150)	
Constant	0.190*** (0.042)	0.214*** (0.060)	54.039*** (5.826)
Observations	140	140	140
R^2	0.186	0.189	0.079
Root MSE	0.440	0.442	51.883

OLS estimates, robust standard errors in parentheses. Observations include the median normalized belief / the median response time of all subjects in the individual treatments. Response time in seconds. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Determinants of correlation neglect

	Median normalized belief			
	(1)	(2)	(3)	(4)
Final high school grade [†]	-0.251*** (0.088)		-0.275*** (0.085)	-0.278*** (0.098)
Raven score [‡]		-0.085** (0.036)	-0.095** (0.036)	-0.095** (0.037)
Constant	1.545*** (0.345)	1.086*** (0.215)	2.171*** (0.582)	2.277*** (0.701)
Additional controls	No	No	No	Yes
Observations	70	70	70	70
R^2	0.094	0.075	0.187	0.211
Root MSE	0.495	0.500	0.472	0.487

OLS estimates, robust standard errors in parentheses. Observations include all median normalized beliefs from the correlated treatment. Additional controls include gender, age, risk aversion, income, working hours, and a high-stakes dummy.

[†] Scale: 1 (worst) - 5 (best). In the German system, the final high school grade (“Abitur”) is a summary statistic of grades in the final years of secondary education and serves as primary university entrance criterion.

[‡] Scale: 0 (worst) - 10 (best). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Using only the sample of subjects in the correlated treatment, Table 4 reports OLS regressions of the median of each individual’s normalized beliefs on our questionnaire measures. The table shows that both better high school grades and higher scores on a Raven IQ test are associated with lower (normalized) beliefs. Thus, both of our proxies indicate that higher cognitive ability leads to more rational beliefs. In combination, these results suggest that neglecting correlations is not the outcome

of a “rational” decision to economize on limited cognitive resources, but that the bias is instead driven by an inability to adequately infer from correlations.

Recently, Oechssler, Roeder and Schmitz (2009) showed that several biases in probability judgment are associated with low cognitive ability. Our results confirm this finding in the context of correlated information. Notably, this result obtains despite the rather homogenous student sample. This suggests that correlation neglect and the resulting belief variance might be even more pronounced in contexts in which people’s cognitive abilities are more heterogeneous.

3.4.2 Market Treatments - Bubbles and Crashes

Price Levels Across Treatments

In both market treatments we have observations from 18 market groups who traded in ten trading rounds each. For each market group and trading round, we define the price of the last concluded trade to be the market price. All results are robust to defining the market price as median or mean trading price. We first consider the effects of our treatment variation on price levels.

Result 4. *Market prices differ between treatments as predicted by correlation neglect. In the correlated treatment, we observe frequent price bubbles or crashes, respectively, depending on the base signal.*

Table 5 provides summary statistics and summarizes key findings for all ten trading rounds. We present the theoretical market price predictions (consisting of the Bayesian and the full correlation neglect belief, respectively) as well as actual price levels. In all rounds but one, prices significantly differ between treatments in the direction one would expect from a correlation neglect perspective. While market prices in the control treatment follow the Bayesian prediction rather closely, we observe frequent price bubbles or crashes in the correlated treatment. Thus, the magnitude of the base signal relative to the other signals consistently predicts whether assets sell above or below the values from the uncorrelated treatment. The last column of the table indicates whether subjects’ beliefs about the value of the assets (which were elicited before trading started) differed significantly between the treatments. This was the case in all rounds but one. In all of these rounds, the belief bias carries over to aggregate outcomes.

Similar to the individual treatments, we now aggregate observations across rounds to get a more robust view on the implications of correlation neglect. Analogously to the belief normalization, we normalize the market price of each round and market group such that it takes on the value zero if it equals the Bayesian belief and

Table 5: Market prices by trading round

True State	Bayesian Belief	Correlation Neglect Belief	Median Market Price Control Treatment	Median Market Price Correlated Treatment	Ranksum Test (p-value)	Beliefs Differ?
10	7.75	9.88	8.35	9.05	0.0093	Yes
88	71.25	96.63	86.5	93.45	0.0338	Yes
250	259.75	219.38	275	260	0.0113	Yes
732	853.15	709.13	820	737	0.0001	Yes
1'000	974.75	1'042.38	1'000	1'039	0.0723	Yes
4'698	4'810	3'209	5'200	4'470.5	0.0085	Yes
7'338	8'604.5	9'277.25	9'124	8'999	0.6087	No
10'000	7'232.25	4'887.63	7'575	6'250	0.0534	Yes
23'112	26'331	20'745.5	24'100	21'300	0.0007	Yes
46'422	38'910.5	25'625	41'000	35'000	0.0015	Yes

Median market prices are defined as the median of all market prices over the 18 markets in the respective round. Beliefs are said to differ between treatments in a particular round if and only if p-value < 0.05 , Wilcoxon ranksum test. Note that subjects faced the ten rounds in randomized order.

the value one if it equals the full correlation neglect belief. Then, for each market group, we use the median normalized market price over the ten rounds for further analysis.

Result 5. *Aggregating the data over the ten trading rounds, we find strong evidence that correlation neglect predictably affects aggregate outcomes.*

Figure 3 provides kernel density estimates of the aggregated data, i.e., of the median normalized prices. It reveals a pronounced and statistically significant difference between the two treatment groups (p-value < 0.0001 , Wilcoxon ranksum test). Normalized prices in the control treatment are centered close to zero, confirming the standard result that double-auctions tend to produce price levels close to fundamentals. Prices in the correlated treatment, however, are centered around 0.6, i.e., prices systematically overshoot in the direction predicted by correlation neglect. Appendix A3 shows that these results neither hinge on the particular aggregation procedure we employ nor on our definition of the market price.

Figure 3 shows that median normalized market prices in the correlated treatment are systematically distorted. However, by construction, this normalization does not allow us to distinguish the occurrence of bubbles from that of crashes. Thus, in what follows, we slightly reformulate our normalization. To be specific, in rounds in which correlation neglect predicts overoptimism, the normalization remains the same, so that a price of one (zero) indicates full correlation neglect (Bayesian price levels). On the other hand, in periods in which neglecting correlations leads to overpessimism, we normalize prices such that full correlation neglect is indicated




Figure 3: Kernel density estimates of median market prices

by (-1) and the Bayesian benchmark by zero, respectively.¹⁸ This normalization makes all rounds comparable, yet still allows for an identification of the direction of potential price distortions.

The two panels in Figure 4 plot the median market prices against their theoretical predictions across our ten trading rounds (we used two different orderings of rounds). First note that by construction the Bayesian price is always given by zero. The full correlation neglect prediction, on the other hand, alternates between one and (-1) depending on whether correlation neglect implies overoptimism or -pessimism. The plots show that in many periods the median price in the correlated treatment follows the correlation neglect prediction, so that prices frequently overshoot as compared to the uncorrelated treatment. Thus, bubbles and crashes predictably alternate across rounds. As a result, the excessive belief swings implied by correlation neglect directly translate into volatile price levels.




Figure 4: Median market prices across trading rounds for the two randomizations

¹⁸Formally, the new set of normalized prices z_i^j is given by $z_i^j = y_i^j \times (2 \times \mathbb{1}_{b_{CN}^j > b_i^j} - 1)$. See footnote 16.

Individual Trading Behavior

So far, we have shown that correlated information structures have predictable consequences for market outcomes, i.e., price levels. Next, we demonstrate that correlation neglect also predicts individuals' trading behavior.

Result 6. *Correlation neglect is reflected in trading behavior. In rounds in which ignoring correlations predicts an upward (downward) biased belief, subjects with a higher propensity to overlook correlations hold significantly more (less) assets. Consequently, these subjects earn lower profits.*

Correlation neglect makes clear predictions about who should hold the assets in the correlated treatment. In rounds where correlation neglect leads to an overvaluation of assets, subjects who ignore correlations should own most of the assets. Likewise, in rounds where correlation neglect leads to an undervaluation of assets, subjects who correctly process the correlation should hold the majority of the assets. To examine this prediction, we split the ten rounds into two groups, namely into rounds in which correlation neglect predicts overly optimistic (pessimistic) beliefs and hence tends to produce a bubble (crash). We then predict asset holdings using the beliefs about the true asset value that subjects stated before trading started. As in our analysis of the determinants of correlation neglect in terms of subjects' characteristics, we employ median normalized beliefs for our analysis.

Columns (1)-(4) of Table 6 provide OLS estimates of the determinants of an individual's asset holdings in the correlated treatment, clustered at the market group level. The regressions establish that the magnitude of the belief bias predicts asset holdings. Columns (1) and (2) show that in trading rounds in which correlation neglect leads to an overly optimistic belief, those subjects with a higher propensity to ignore correlations hold significantly more assets. Likewise, in rounds in which the bias implies overpessimism, those subjects whose stated beliefs reveal a higher degree of correlation neglect hold fewer assets (columns (3) and (4)). These effects are both statistically and economically significant. For example, during crashes, the point estimate suggests that moving from a Bayesian to a full correlation neglect belief implies a reduction in asset holdings of 1.3, which is almost a third of each seller's initial endowment.

Columns (5) and (6) investigate the implications of correlation neglect for subjects' trading profits. The results show that the belief bias and the corresponding trading behavior described above are associated with lower profits. Also recall that we established a close connection between cognitive ability and correlation neglect. In combination, our results thus suggest that especially people with low cognitive skills are likely to make systematically distorted choices in market interactions.

Table 6: Determinants of asset holdings and profits in the correlated market treatment

	Median asset holdings in crashes		Median asset holdings in bubbles		Median profit	
	(1)	(2)	(3)	(4)	(5)	(6)
Median correlation neglect	-1.534*** (0.174)	-1.300*** (0.187)	0.642*** (0.118)	0.259* (0.120)	-0.014** (0.006)	-0.013** (0.005)
Constant	2.853*** (0.123)	2.460** (1.049)	1.432*** (0.135)	3.821*** (1.201)	1.008*** (0.004)	1.018*** (0.027)
Additional controls	No	Yes	No	Yes	No	Yes
Observations	143	143	143	143	143	143
R^2	0.306	0.425	0.200	0.422	0.051	0.160
Root MSE	1.158	1.077	1.293	1.123	0.023	0.022

OLS estimates, standard errors clustered at the market level. Observations include median asset holdings / profits of all subjects in the correlated treatment. Bubbles (crashes) are defined as rounds in which correlation neglect predicts overoptimism (-pessimism). The median profit is computed as the median normalized profit from all rounds, where for each trader and for each round a normalized profit is defined as $\pi = \frac{\text{Money holdings} + \text{value of assets held}}{\text{Monetary value of endowment}}$, where for sellers (buyers) the value of the endowment consists of the value of the initially owned assets (the budget). Median correlation neglect in (1) and (2) [(3) and (4)] is computed as the median normalized belief of the rounds in which correlation neglect predicts overoptimism [overpessimism]. In (5) and (6), median correlation neglect equals the median normalized belief across all rounds. Additional controls include a buyer dummy, age, gender, income, high school grade, and working hours. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Markets and Individual Biases

Thus far, we have established that subjects' tendency to overlook correlations in information sources produces distorted aggregate outcomes. Now, we analyze to what extent the market we implement and the information that is provided through the market mitigate correlation neglect.

Result 7. *We do not find evidence indicating that the introduction of the particular competitive market environment we implement reduces the bias caused by neglecting correlations.*

We present two pieces of evidence to support this claim. First, we ask if the bias reflected in market prices is (on average) smaller than the individual belief bias, as one would expect from an attenuation effect. Using a difference-in-difference OLS estimation, Appendix A4 shows that this is not the case. Rather, the degree of correlation neglect is very similar between the individual and the market treatments.

Second, in our market setup, subjects could learn by observing others as well as through the feedback provided at the end of each trading round. If learning played an important role, then the price distortion should be reduced towards the end of the experiment. However, we find no evidence for such an effect - neither beliefs nor prices show any sign of converging to their counterparts in the uncorrelated treatment. For instance, if we take the last round from all market groups and normalize the market price (to make it comparable between different orderings of

rounds), we still find a significant treatment difference (p-value = 0.0290, Wilcoxon ranksum test). In an OLS regression of all normalized market prices on a linear and a quadratic time trend, the time regressors are jointly insignificant, see Appendix A2 for details. Thus, we find no evidence supporting the learning arguments.

In sum, our results suggest that at least in the environment we implement, market forces are not sufficient to mitigate correlation neglect.

4 Neglect of Correlations in Newspaper Articles

In our baseline experiment, we deliberately designed an abstract decision environment which allowed tight control over signal precision as well as the degree of correlation. To show the robustness of our findings, we now make use of a naturally occurring correlation in an informational context familiar to subjects.

Experimental Design: In the experiment, subjects estimated the growth of the German economy in 2012. Subjects were asked to read (shortened) real newspaper articles discussing and summarizing growth forecasts and then provided an incentivized estimate. Again, we study two main treatments, one in which information is correlated and one in which it is not. In the correlated treatment, subjects received two articles. One of these articles discussed a joint forecast from April 2012. The joint forecast is determined on a biannual basis in a cooperation of several German research institutes, thus aggregating information of the participating institutes. It predicted that the German economy would grow at a rate of 0.9 % in 2012. The other article discussed a forecast of one particular institute from March 2012 that predicted a growth rate of 1.3 %. However, since this institute also takes part in the joint forecast, the information from that institute is already incorporated in the joint forecast. In other words, the two articles are correlated since information from the second article is captured in the article presenting the joint forecast. This correlation was salient, since the article reporting the joint forecast clearly stated all participating institutes.¹⁹

As a control condition, we merely supplied the joint forecast. Since the joint forecast reflects the information contained in the second article, this treatment removes the correlation, yet keeps the informational content identical.

In this newspaper context, detecting the correlation might be an additional challenge as compared to processing it. Thus, we also conducted an additional treatment in which we increased the salience of the correlation. The treatment was identical to the correlated treatment except for an additional summary at the end of the two

¹⁹See Appendix B for a translation of the articles.

articles. This summary again stated that the institute predicting a growth rate of 1.3 % also participated in the joint forecast.²⁰

Procedural Details: 284 subjects participated in the experiment. Sessions were conducted using paper and pencil in the BonnEconLab at the end of different and unrelated experiments. Treatments were randomized within session. In the conditions involving two articles, the order of the articles was randomized. The study took five minutes on average. At the end of each session, one subject was randomly selected for payment. He was asked to write his address on an envelope and was reminded that his earnings will be sent to him as soon as the official growth figures are available. Earnings were 10 euros if the estimate turned out to be correct. For every 0.1 percentage point deviation, 1 euro was deducted. Negative earnings were not possible. The randomly selected subjects earned 7.30 euros on average.

Prediction: In the correlated condition, subjects only received information which was already included in the joint forecast. However, if these subjects fail to take the correlation between the two articles into account, average estimates would be higher than those in the control treatment. If part of the problem was to detect the correlation (rather than to process it), we would expect less correlation neglect in the treatment with increased salience.

²⁰There are four potential concerns with respect to our design. First, one could argue that the difference between the joint forecast of 0.9 % and the forecast of 1.3 % is informative because it indicates a high variance of forecasts. This variance in turn might allow inference about the signal precision of the participating institutes. Consequently, subjects in the correlated condition could put lower weight on the forecasts (relative to their own prior) when determining their estimate. Notice, however, that even if subjects actually went through this kind of inference, this would not explain our treatment difference. Findings from our control condition reveal that subjects' priors are on average actually slightly below the joint forecast of 0.9 %. Thus, lower weight on the joint forecast in the updating process would not lead to estimates that are closer to 1.3 %. A further potential concern might be that information from the second article is informative if the forecast of the institute that is discussed in this article is not appropriately incorporated in the joint forecast. This does not seem plausible. However, to further address this issue, we asked a subset of subjects at the end of the experiment if they had the suspicion that this is actually the case. In the main correlated treatment, only two subjects (7 %) indicated such a concern. Our findings remain qualitatively unchanged if we remove these subjects from the subsample. Third, subjects could interpret the mere presentation of the article discussing the forecast of 1.3 % as an indication that the article has to be of informational value. We addressed this concern by introducing a new treatment, which is identical to the correlated treatment except that it contains a second incentivized question which relates to labor market information provided in the article discussing the 1.3 % forecast. Thus, there was a natural reason for the presence of the second article, which was unrelated to the question about GDP growth. The results suggest that this type of effect does not drive our results. Estimates in this treatment are almost identical to those in the standard correlated condition and significantly different from those in the control condition (p-value < 0.0001, Wilcoxon ranksum test). Finally, the two forecasts were published one month apart from each other. This is unproblematic, however, since the joint forecast was released at the later date. Thus, the timing as such provided no reason for subjects to place any weight on the 1.3 % forecast.

Results: The median estimate over all treatments was a growth rate of 1.035 %. Focusing on the control condition, the median estimate was 0.82 %, slightly more pessimistic than the joint forecast of 0.9 %. Only about 10 % of subjects in the control treatment estimated a growth rate of exactly 0.9 %, suggesting that many subjects also used prior information for their estimate. In the correlated treatment, the median estimate was 0.28 percentage points higher as compared to the control treatment, a statistically significant difference (p-value < 0.0001, Wilcoxon ranksum test). Thus, subjects neglect the correlation between the two newspaper articles, i.e., the two forecasts, leading to substantial overoptimism regarding GDP growth.

In the additional treatment which increased the salience of the correlation, we still find significant evidence for correlation neglect. The median estimate in this treatment is also 1.1 % and estimates differ significantly from those in the control treatment (p-value < 0.0001, Wilcoxon ranksum test). Thus, the results suggest that in the very simple environment we implement, detecting correlations was not a first-order problem as compared to processing them.

In sum, these findings emphasize the robustness of correlation neglect in belief formation. They show that people end up with distorted beliefs even in familiar environments such as reading and learning from newspaper articles. While subjects may not frequently be required to predict GDP growth as such, the type of information we provide is very typical for everyday information processing.

5 Concluding Remarks

This paper provides evidence for people’s tendency to neglect correlations in information sources when forming beliefs. We show that this bias has implications both at an individual and at the market level. Specifically, neglecting correlations causes overshooting beliefs which predictably drive price bubbles and crashes. People’s problems to understand correlations appear to be robust to variations in the decision making context. Neither an increase in the stake size, nor market interaction, nor the introduction of a familiar informational environment eliminated the cognitive bias. In concluding the paper, we briefly comment on the origins and potential aggregate implications of correlation neglect.

While bounded rationality in a broad sense provides an intuitive rationale for our findings, an interesting question is which precise mechanism causes the bias. Our results suggest that neglecting correlations is not the outcome of a “rational” decision to sacrifice correct beliefs for the sake of economizing on scarce cognitive resources. Rather, correlation neglect seems to be driven primarily by a general in-

capability to understand the implications of statistical correlations.²¹ This interpretation conforms well with Kahneman and Tversky (1973) who note that correlated information sources tend to produce consistent signals and (due to the so-called representativeness heuristic) might hence lead to an “illusion of validity”. These arguments suggest that teaching people the statistical skills needed to infer from correlated signals might prove useful. In this respect, an interesting question is how much it would take to teach people these skills. In our main experiment, while the information structure was clear, simple, and salient, we did not give subjects a “recipe” on how to extract the underlying independent signals from the information they received. It would be informative to learn whether this would actually be needed or whether there is an intermediate step at which subjects already fully grasp the required computations.

In our main experiment, we made sure that subjects were aware of the correlation between the different sources of information. However, an additional challenge in many field situations might be to detect correlations in the first place. The distinction between failure to notice and failure to process correlations is important, as it points to two different approaches to help people reduce the bias caused by neglecting correlations. In our newspaper study, discovering correlations did not seem to be a major problem as increasing the salience of the correlation had no effect on the magnitude of the bias. However, in more complex situations, correlations might well be less salient and much more difficult to detect.

As for the aggregate implications of correlation neglect, our results bear an interesting relationship to the voluminous literature on bubbles, crashes, and excess price volatility.²² In this respect, our findings could be viewed as lending micro-founded support to the familiar claim that sometimes markets tend to undergo excessive confidence swings. In particular, in the ubiquitous presence of correlated information - through the telling and retelling of similar stories - aggregate distortions might well have their origins in systematically biased beliefs.

²¹One possible reason for this is that people could intuitively focus their attention on actual signal realizations as opposed to the structure underlying the signals. See Kőszegi and Szeidl (2013) and Bordalo, Gennaioli and Shleifer (2012) for theories of attention and focusing.

²²There is a large theoretical and experimental literature on rational as well as irrational bubbles and crashes. See e.g. Smith, Suchanek and Williams (1988) for seminal experimental work and Brunnermeier and Oehmke (2012) for a mostly theoretical survey of the literature.

References

- Akerlof, G.A. and R.J. Shiller**, *Animal spirits: How human psychology drives the economy, and why it matters for global capitalism*, Princeton University Press, 2009.
- Bénabou, R.**, “Groupthink: Collective delusions in organizations and markets,” *Review of Economic Studies*, 2013, 80 (2), 429–462.
- Bordalo, P., N. Gennaioli, and A. Shleifer**, “Salience theory of choice under risk,” *Quarterly Journal of Economics*, 2012, 127 (3), 1243–1285.
- Brunnermeier, M. and M. Oehmke**, “Bubbles, financial crises, and systemic risk,” *Handbook of the Economics of Finance*, 2012, 2.
- Brunnermeier, M.K. and J.A. Parker**, “Optimal expectations,” *American Economic Review*, 2005, 95 (4), 1092–1118.
- Camerer, C.F.**, “Do biases in probability judgment matter in markets? Experimental evidence,” *American Economic Review*, 1987, 77 (5), 981–997.
- Corazzini, L., F. Pavesi, B. Petrovich, and L. Stanca**, “Influential listeners: An experiment on persuasion bias in social networks,” *European Economic Review*, 2012, 56, 1276–1288.
- DeGroot, M.H.**, “Reaching a consensus,” *Journal of the American Statistical Association*, 1974, 69 (345), 118–121.
- DellaVigna, S. and M. Gentzkow**, “Persuasion: Empirical evidence,” *Annual Review of Economics*, 2010, 2 (1), 643–669.
- DeMarzo, P.M., D. Vayanos, and J. Zwiebel**, “Persuasion bias, social influence, and unidimensional opinions,” *Quarterly Journal of Economics*, 2003, 118 (3), 909–968.
- Eyster, E. and G. Weizsäcker**, “Correlation neglect in financial decision-making,” *Discussion Papers of DIW Berlin*, 2011.
- **and M. Rabin**, “Naïve herding in rich-information settings,” *American Economic Journal: Microeconomics*, 2010, 2 (4), 221–243.
- Fischbacher, U.**, “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, 2007, 10 (2), 171–178.

- Forsythe, R., F. Nelson, G.R. Neumann, and J. Wright**, “Anatomy of an experimental political stock market,” *American Economic Review*, 1992, *82* (5), 1142–1161.
- Ganguly, A.R., J.H. Kagel, and D.V. Moser**, “Do asset market prices reflect traders’ judgment biases?,” *Journal of Risk and Uncertainty*, 2000, *20* (3), 219–245.
- Gneezy, U., A. Kapteyn, and J. Potters**, “Evaluation periods and asset prices in a market experiment,” *Journal of Finance*, 2003, *58* (2), 821–838.
- Golub, B. and M.O. Jackson**, “Naïve learning in social networks and the wisdom of crowds,” *American Economic Journal: Microeconomics*, 2010, *2* (1), 112–149.
- Grebe, T., J. Schmid, and A. Stiehler**, “Do individuals recognize cascade behavior of others? – An experimental study,” *Journal of Economic Psychology*, 2008, *29* (2), 197–209.
- Greiner, B.**, “The online recruitment system ORSEE 2.0 – A guide for the organization of experiments in economics,” *University of Cologne, Working paper series in economics*, 2004, *10* (23), 63–104.
- Kahneman, D. and A. Tversky**, “On the psychology of prediction,” *Psychological Review*, 1973, *80* (4), 237–251.
- Kallir, I. and D. Sonsino**, “The neglect of correlation in allocation decisions,” *Southern Economic Journal*, 2010, *75* (4), 1045–1066.
- Kluger, B.D. and S.B. Wyatt**, “Are judgment errors reflected in market prices and allocations? Experimental evidence based on the Monty Hall problem,” *Journal of Finance*, 2004, *59* (3), 969–998.
- Kőszegi, Botond and Adam Szeidl**, “A model of focusing in economic choice,” *Quarterly Journal of Economics*, 2013, *128* (1), 53–104.
- Kübler, D. and G. Weizsäcker**, “Are longer cascades more stable?,” *Journal of the European Economic Association*, 2005, *3* (2-3), 330–339.
- Oechssler, J., A. Roider, and P.W. Schmitz**, “Cognitive abilities and behavioral biases,” *Journal of Economic Behavior & Organization*, 2009, *72* (1), 147–152.
- Penczynski, S.**, “A level-k model of social learning,” *Working Paper*, 2012.

- Selten, R.**, “Axiomatic characterization of the quadratic scoring rule,” *Experimental Economics*, 1998, 1 (1), 43–62.
- Smith, V.L., G.L. Suchanek, and A.W. Williams**, “Bubbles, crashes, and endogenous expectations in experimental spot asset markets,” *Econometrica*, 1988, 56 (5), 1119–1151.
- Sobel, J.**, “On the relationship between individual and group decisions,” *Working Paper*, 2012.
- Stoner, James A.F.**, “Risky and cautious shifts in group decisions: The influence of widely held values,” *Journal of Experimental Social Psychology*, 1968, 4 (4), 442–459.

FOR ONLINE PUBLICATION

Appendix A

A1 The Conceptual Framework with N Signals

Suppose now that nature generates N independent signals $s_i = \theta + \epsilon$, $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. The decision maker observes s_1 and

$$\tilde{s}_i = \alpha s_1 + (1 - \alpha)s_i, \quad i \in \{2, \dots, N\}$$

As above, the decision maker's inference of the underlying independent signal is given by $\hat{s}_i = s_i + \chi\alpha(s_1 - s_i)$. We then define an agent's bias as the absolute expected difference between a (partially naïve) mean belief and the mean Bayesian posterior conditional on some base signal $s_1 \neq \theta$,

$$abs(E[m_{CN} - m_B | s_1]) = abs \left(E \left[\frac{s_1 + \sum_{i=2}^N \hat{s}_i}{N} - \frac{s_1 + \sum_{i=2}^N s_i}{N} \middle| s_1 \right] \right)$$

where the expectation is taken with respect to the true data-generating process. This simplifies to

$$\begin{aligned} abs(E[m_{CN} - m_B | s_1]) &= abs \left(E \left[\frac{\chi\alpha}{N} \sum_{i=2}^N (s_1 - s_i) \middle| s_1 \right] \right) \\ &= abs(s_1 - \theta) \frac{N-1}{N} \chi\alpha \end{aligned}$$

which is a strictly increasing function of N as long as $\chi\alpha > 0$, i.e., as long as there is positive correlation and the agent is marginally naïve. Thus, even in information-rich environments ($N \rightarrow \infty$), agents come to hold biased beliefs.

A2 (No) Learning Over Time

Columns (1) and (2) of Table 7 provide the results of OLS regressions of all normalized beliefs in the individual correlated treatment on a linear and a quadratic time trend. They show that normalized beliefs do not become smaller over time, i.e., they do not converge to the Bayesian belief of zero. In columns (3) and (4), we show that beliefs do not converge to their counterparts in the uncorrelated treatment either. To this end, we take all normalized beliefs from the correlated treatment, subtract the median normalized belief of the median individual in the respective trading round in the uncorrelated treatment and then regress this modified belief on a time trend (in effect, this accounts for fixed effects of specific trading rounds). The results show that the difference between the correlated and the uncorrelated treatment does not become smaller over time.

Similarly, Table 8 gives an overview of the time trend of market prices. In columns (1) and (2), we report the results of OLS regressions of all normalized market prices in the correlated treatment on a linear and a quadratic trend. The time regressors are individually and jointly insignificant, indicating that market prices do not converge to Bayesian levels. We also show that prices do not converge to their counterparts in the uncorrelated treatment (columns (3) and (4)). To this end, we take all normalized market prices and then subtract the median normalized market price of the median market group in that round in the uncorrelated treatment.

Table 7: Time trend of beliefs in the correlated treatment

	Normalized belief		Normalized belief minus median belief in uncorrelated	
	(1)	(2)	(3)	(4)
# of round	-0.001 (0.016)	0.119 (0.101)	-0.007 (0.015)	0.107 (0.102)
# of round squared		-0.011 (0.009)		-0.010 (0.009)
Constant	0.715*** (0.128)	0.477** (0.200)	0.685*** (0.127)	0.458** (0.201)
Observations	693	693	693	693
R^2	0.000	0.002	0.000	0.002
p-value of F-statistic	0.957	0.480	0.645	0.454

OLS regressions, standard errors (clustered at individual) in parentheses.

Observations include all normalized beliefs from all rounds in the correlated individual treatment. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Time trend of market prices in the correlated treatment

	Normalized market price		Normalized market price minus median price in uncorrelated	
	(1)	(2)	(3)	(4)
# of trading period	-0.005 (0.087)	0.451 (0.484)	-0.015 (0.083)	0.489 (0.485)
# of trading period squared		-0.042 (0.038)		-0.046 (0.038)
Constant	0.850*** (0.212)	-0.056 (0.850)	0.728 (0.178)	-0.274 (0.841)
Observations	170	170	170	170
R^2	0.000	0.004	0.000	0.005
p-value of F-statistic	0.9552	0.2646	0.856	0.1477

OLS regressions, standard errors (clustered at market group level) in parentheses. Observations include the market prices from all trading rounds in the correlated market treatment, excluding one extreme outlier (normalized price = 1398). All results are robust to including this observation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A3 Robustness of Results in Market Treatments

This appendix provides a robustness check for our main treatment effect in the market treatments. In particular, we show that neither our definition of the market price nor our aggregation procedure drives our results. Using three definitions of market prices and two different aggregation procedures (for aggregating the market prices of ten trading rounds into a single price per market group), Table 9 presents the p-value of ranksum tests for the equality of the aggregated market price between treatments.

Table 9: P-values of Wilcoxon ranksum tests for equality of aggregated market price between treatments.

Aggregation mechanism \ Market price	Median price	Average price	Last trading price
Median market price	0.0000	0.0000	0.0000
Average market price	0.0001	0.0002	0.0054

A4 Comparison Between (Correlated) Individual and Market Treatments

This section asks whether the bias reflected in market prices is smaller than the bias inherent in individual beliefs in the individual belief elicitation treatments. For this purpose, we pool all median normalized beliefs from the individual treatments with all median normalized prices from the market treatments. We thus end up with 140 normalized beliefs (one per individual) and 36 normalized market prices (one per market group), i.e., with 176 observations. Table 10 then presents the results of a difference-in-difference type OLS regression, in which we regress correlation neglect on a treatment group dummy, a market dummy, and an interaction term equal to one if an observation stems from the correlated market treatment. If the correlation neglect bias reflected by market prices is smaller than the individual belief bias, then this interaction term should have a negative coefficient. The fact that this is not the case, however, suggests that in our environment the outcome of the market interaction (the price) is not less biased than individual beliefs.

Table 10: Market and individual treatments

	Median correlation neglect
1 if correlated	0.428*** (0.108)
1 if market	-0.064 (0.073)
1 if correlated market	-0.004 (0.129)
Constant	0.214*** (0.060)
Observations	176
R^2	0.224
Root MSE	0.405

OLS estimates, robust standard errors in parentheses.

Observations include all median normalized beliefs from the individual treatments and all median normalized prices.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Appendix B

Newspaper Articles

Translated into English

Manager-Magazin, 14.03.2012

IWH increases growth forecast

The German economy seems to be gaining speed. According to the Institute for Economic Research Halle, the short period of economic weakness is over. Thus, the researchers increase their growth forecast for Germany significantly.

On Wednesday, the institute in Halle announced that it expects the German economy to grow by 1.3 % this year. According to the IWH experts, the risks relating to the debt and trust crisis in Europe have been slightly reduced. Both the world economy and the German economy are said to have started significantly better into 2012 than was projected in autumn 2011. According to the IWH, the positive economic development will also affect the labor market.

Welt Online, 19.04.2012

Leading economic research institutes say German economy is in upswing

According to leading economic research institutes, the German economy is in upswing. In their joint “Spring 2012” forecast, published on Thursday, the institutes forecast a growth of the German economy of 0.9 %.

According to the researchers, the biggest “down-side risk” for the future remains to be the debt and trust crisis in the Euro area. While the remarkable measures of the European Central Bank relieved stress in the banking system, they are not more than a gain of time.

The forecast is prepared by the Ifo Institute in Munich, the ETH Zurich, the ZEW Mannheim, the Institute for Economic Research Halle, Kiel Economics, IHS Vienna, and the RWI Institute in Essen.