Is Overconfidence a Judgment Bias? Theory and Evidence*

Stephen V. Burks[1], Jeffrey P. Carpenter[2], Lorenz Goette[3], Aldo Rustichini[4]

June 12, 2009

Abstract

Accurate assessments of one's abilities are important in many domains. These assessments, e.g, form the basis for the choice of education, or the wage one is asking in wage negotiations. Evidence from psychology and economics indicates that many individuals seem to overestimate their ability. These results are typically taken as evidence that individuals are overconfident in their abilities.

Recent modeling efforts in economics have shown that this need not be irrational, but can be the outcome of rational Bayesian updating, thus providing a potential explanation for overconfidence (Benoît and Dubra, 2007).

In this paper, we provide three results: First, we show that Bayesian updating imposes restrictions on ability judgments that can be tested empirically. Using data on 1015 individuals' confidence judgments about two cognitive tests, we test these restrictions and reject the hypothesis that rational updating explains overconfidence. Second, we test whether self-image concerns contributes to overconfidence (Kőszegi, 2006; Weinberg, 2006). Our results suggest that individuals do not care about inferences created by signals of their ability, but rather consume the signals themselves - they like to hear that they are good. Third, we provide evidence that personality characteristics strongly affect confidence judgments: Socially dominant individuals have more confident judgments about their ability. These judgments are unfounded, as socially dominant individuals do not have higher ability. Thus, our evidence suggests that overconfidence is a judgment bias, and that social dominance plays an important role in explaining it.

[1] Division of Social Sciences, University of Minnesota, Morris, 600 East 4th Street, Morris, MN 56267-2134

[2] Department of Economics, Middlebury College, Middlebury, VT 05753

[3] Department of Economics University of Geneva, 40 Bd. du Pont d'Arve, 1217 Geneva, Switzerland

[4] Department of Economics, University of Minnesota, 1925 4th Street South 4-101, Hanson Hall, Minneapolis, MN 55455-0462

*The authors gratefully acknowledge generous financial support for the Truckers and Turnover Project from the John D. and Catherine T. MacArthur Foundation's Research Network on the Nature and Origin of Preferences, the Alfred P. Sloan Foundation, the Trucking Industry Program at the Georgia Institute of Technology, the University of Minnesota, Morris, the Federal Reserve Bank of Boston and the University of Geneva, and financial and in-kind support from the cooperating motor carrier, its staff, and its executives. The views expressed are those of the authors, and do not necessarily reflect the views of any neither of their employers nor of the project's funders.

1 Introduction

Well-calibrated judgments about one's abilities are important in many economic decisions. However, there is evidence from studies in psychology and economics suggesting that individuals have too much confidence in their abilities. In a typical study, almost nobody rates themselves in the bottom 40 percent of a distribution, largely independent of the skill in question. This is such a robust finding that De Bondt and Thaler (1985) call it "perhaps the best-documented anomaly" (or something like that). Studies also link measures of what the authors argue is overconfidence to behaviors, and show that more confident judgments are associated with more daring behaviors. Malmendier and Tate (2008) show that more confident CEOs make more daring merger decisions (see also Malmendier and Tate, 2005). Dohmen and Falk (2006) and Bartling and Fehr (2009) show that overconfident individuals are more likely to select themselves into competitive payment schemes. Barber and Odean (2001) show that men engage in more frequent trading in common stock, consistent with the evidence from psychology that men are more overconfident than women. The trading reduces their returns substantially relative to women. Thus, if overconfidence is truly a judgment bias, these studies should raise concern, as they raise the possibility that individuals act on biased beliefs.

If individuals had perfect knowledge of their abilities, results showing that, e.g., 50 percent of the individuals rate themselves in the top 25 percent of an ability distribution necessarily imply a judgment bias.[1] However, assuming perfect knowledge of one's ability may not be realistic. Rather, individuals may only vaguely know their abilities, and update their beliefs as new information arrives. A recent paper by Benoît and Dubra (2007) shows that, if individuals have imperfect knowledge of their own ability, even perfectly rational Bayesian updaters may report overconfident beliefs in a typical study.

Benoît and Dubra point out that, in most studies, individuals are asked to indicate their most likely place in the ability distribution. They provide a general characterization of the information structure leading to results that, for example, 50 percent of the individuals put themselves in the top 25 percent of the ability distribution. Intuitively, this can arise if the signals individuals receive become more noisy the better the signals are, akin to taking an easy test: Everyone who fails the test can be sure that his ability is low. However, low-ability types sometimes also pass the test by sheer luck. But still, passing the test rationally leads individuals to believe it is more likely that they have high ability, therefore creating 'overconfidence' by this measure. Several papers (Kőszegi, 2006; Weinberg, 2006) also provide plausible psychological underpinnings, showing that such types of information structures can arise endogenously. They argue that image concerns by individuals lead individuals with high beliefs to refrain from seeking more information, leading to an information structure that is conducive to creating overconfidence. Yet, while overconfidence in that sense may prevail

[1] Merkle and Weber (2009) do show overconfidence leads to bias in beliefs. Their test is based on eliciting the c.d.f. of beliefs over abilities for which it is difficult to pin down the true distribution. This allows them to reject Bayesian updating without even knowing what the true distribution of ability is.

in the population, the beliefs are still unbiased, as the individuals who think they are in the upper half of the distribution recognize that this is not sure and, in their model, attach the correct probability to this state. As a consequence, individuals may take actions based on these beliefs, which may not be optimal compared to the case in which the individual knew his true type. However, conditional on the information that the individual received, the judgments are consistent and their actions optimal.

In this paper, we show that if overconfident judgments are the result of Bayesian information processing from a common prior, this places testable restrictions on the beliefs as a function of the individual's true ability. As we show, it must be true that of all individuals placing themselves in ability quantile k, individuals of ability quantile k must be most likely to do so. Therefore, we can base a test of Bayesian overconfidence on whether this is the case. We test the model with data from 1012 subjects, judging their ability for each of two cognitive tests that we administer to them. We clearly reject the restriction: In general, individuals from an ability quantile $j < k$ are more likely to think they are in quantile k.

Out test evaluates the joint hypothesis of Bayesian updating and the common prior assumption, leaving the question unanswered which part of the joint hypothesis has failed. For the class of model relying on image concerns to generate overconfident beliefs, we provide an alternative test that is independent of the common prior assumption. In models of image concerns, individuals would like to belief that they have high ability, but their beliefs are constrained by Bayesian updating (Kőszegi, 2006; Weinberg, 2006). The core of these models is that once individuals are sufficiently certain that they are of high ability, they stop seeking information, as this only offers the downside risk of erroneously revising their belief to lower ability. By contrast, individuals with a low belief "don't take no for an answer" and seek information as long as there is a chance to find out that they may be better. We test this prediction: We offer our subjects the opportunity to find out exactly how well they did in the tests relative to the other participants. Our data strongly reject the prediction of self-image models. We find, in contrast, that individuals with high beliefs are *more likely* to demand information about their ability.

Thus, beliefs do play an important role in demanding information, but not in ways that is consistent with preserving self-image. Our results are more in line with a model in which individuals enjoy signals, rather than the resulting belief, confirming that they are good. We further corroborate this interpretation by examining how individual personality differences affect confidence judgments. Consistent with our interpretation, we find that more socially dominant individuals make more confident judgments, holding constant their actual ability. This effect is also quantitatively large: Of those individuals with a below-median score in social dominance, only 33 percent think they are in the top 20 percent of the IQ distribution. Of the individuals with an above-median score in social dominance, 55 percent think they are in the top 20 percent, when, in fact in both groups there is no difference in the frequency of being in the top 20 percent. We also find that trait neuroticism reduces confidence by a similar magnitude.

Overall, our findings suggest that overconfidence arises due to biased interpretation of information. Our results show that overconfidence cannot arise from Bayesian updating on signals about one's ability. We also show that individuals with high confidence are more likely to seek information about their ability, which should provide them with information that undermines overconfidence. This strongly suggests that individuals interpret information in a biased way. This view receives further support from the fact that social dominance makes judgments more confident while neuroticism makes judgments less confident. Thus, these two personality traits shift confidence in the way that would be predicted if they biased the interpretation of information.

The remainder of this paper is structured as follows: Section 2 describes our empirical setup and presents the basic findings on confidence judgments. Section 3 introduces a framework of incomplete information about one's own ability, derives restrictions that this places on confidence judgments, and tests them. Section 4 discusses image preservation as a source of overconfidence, and provides an empirical test. Section 5 presents evidence on how personality traits are related to overconfidence.

2 The Empirical Setup

2.1 Design of the Study

The study was part of the Truckers and Turnover project (see Burks et al., 2007, for a detailed description). The subjects participated in two tests of cognitive ability: An IQ test and a Numeracy test. The two tests were part of a larger data collection process. The sequence of events was the same in both tests and is illustrated in Figure TK: Before the test, the subjects were told of the type of test that was about to follow. This information also included all the specific instructions of how to complete the test, which differed between the IQ test. The IQ test was computerized, while the numeracy test was paper-and-pencil. In order to keep the motivation with the task high, we randomly selected one of the subjects after the test and paid him or her two dollars for every correct answer in the numeracy test (or one dollar for every correct answer in the IQ test). The subjects were informed of this in advance.

After the instructions, we recorded the first self-assessment of the subjects' abilities: The subjects were asked how well they thought they would do in this test relative to the rest of the session's participants. The subjects had to rate their performance in quintiles. If their assessment was correct, they were paid two dollars. This simple mechanism makes truthful answering incentive compatible. The subjects then participated in the test. After the test, the subjects were asked again how well they now thought they did. Again, they had to rate their performance in quintiles, and would be paid another two dollars if correct.

Central to our study, we then asked the subjects whether they would like to learn on Monday (the day the subjects got paid for all of the experiments) their exact score and how well they did in comparison to the others. This is our measure of demand for information: Answering yes will produce an additional signal of their ability, answering no will not.

The exact same procedure was used in the two tests. In total, 1063 subjects over 2 experimental sessions each participated in this study. In all sessions, the IQ test was administered before the numeracy test.

This simple procedure provides incentives to the subjects to get the assessment of their abilities right. It also asks the subjects about a specific group of people that they have known for more than a week by the time of the experiment. Therefore, our design also rules out that the subjects were comparing themselves to groups outside the lab, unlike the most common studies of overconfidence in the psychology literature. Further, it avoids the ambiguities of earlier studies that asked individuals whether they were above or below the mean.[2]

2.2 Descriptive Statistics

Table 1 presents the descriptive statistics. The first panel in the table shows the number of correct answers in the two cognitive tests. It is difficult to compare the performance here to other populations, because, for instance, we only performed one part of the numeracy test, and data for only parts of the test are not readily available. In Burks et al. (2007), it is shown, however that our subjects tend to have somewhat lower IQ scores than representative samples. In particular the top end of the distribution seems to be somewhat less populated in our sample.

Turning to the demographics of hour sample, we see that the most frequent education level in our sample is a high school degree, though some have also degrees from technical schools, and a significant fraction has at least some college education. Considering the blue-collar nature of the job at hand, the education levels are slightly above what one would find in typical other blue-collar jobs. The table further shows that our sample is predominantly Caucasian, and predominantly male, and relatively young for blue-collar workers. See Burks et al. (2007, 2009) for a more extensive discussion.

2.3 Descriptive Evidence on Confidence Judgments

In this subsection, we present the basic evidence on confidence judgments in our study. This serves two purposes. First, we show that our results are comparable to confidence judgments found in other studies. Second, it serves to motivate the theoretical model we discuss in the following sections.

Figure 1 displays the distribution of confidence judgments across all individuals. It shows the typical pattern found in a large number of studies: Very few individuals rate their ability

[2] If, e.g., the mean of abilities is significantly below the mean, a large fraction could correctly answer that they are better than average, which makes the interpretation of these studies difficult.

in the bottom 40 percent of the ability distribution. By contrast, well above 60 percent think they are in the top 40 percent. The figure shows a very similar pattern for the confidence judgments in the two tests. These results are typical for the

Figure 2 shows the confidence judgments as a function of the true ability in the IQ test. Panel The figure displays under- and overconfident judgment relative to the true ability of the individual. Shadings indicate the the extent of overconfidence: Light shading indiviates that the individual is just one quintile off, darker shading indicates that the individual is more than one quintile off. Panel A uses the confident judgments before the IQ test. The figure shows that overconfident judgments are pervasive across the ability spectrum, except where mechanically impossible in the top ability quintile. The figure also shows that the confidence judgments are strongly asymmetric: Underconfidence is much rarer than overconfidence. Panel B in Figure 2 uses the confidence judgments after the IQ test and shows essentially the same pattern. Figure 3 shows the confidence judgments from the numeracy test. The results are vitrutally the same.

3 The Bayesian Model

We want to establish in this section the benchmark model of a population that is forming beliefs on their own ranking on the basis of a rational and unbiased process of updating given the information they have available. Our aim is to show that such a model can produce some features of the confidence judgments that we showed in the previous section, but to also derive testable restrictions imposed by the Bayesian theory.

In the model we consider a large population of individuals, each one endowed with a type $t \in \Omega$, which is the value of a specific characteristic. For example the type of an individual can be his height, something easily determined and observed. Another more interesting example is his ability to score in an IQ test, a quality that can be briefly described as the individual's *IQ*. We are interested in types that are ordinal qualities. In what follows, we will restrict attention to judgments about the individual's position in the distribution of outcome. As in our empirical exercise, we elicit judgments about the quintiles, we also restrict our notation in the model to quintiles in order to minimize notational burden.

Types are determined independently, according to a known probability measure on the set of types. Thus, the population has a common prior on the distribution of types, which is the uniform distribution. We assume that the prior is common.

Individuals do not know their type, but during their life they gather information by observing private signals depending on their type. On the basis of this information they update in a Bayesian fashion their belief on their own type, which initially was the common prior, and therefore also they update the belief they have on their own relative position in the population with respect to the characteristic we are considering. For example, through their school performance, job performance, as well as occasional exchanges with other people they form an opinion on their IQ, and hence of their relative standing within the population. Formally, we assume that individuals observe an outcome $x_i \in X, i = 1, ...n$ from the signal structure.[3]

When a subject participates in an experiment like ours he comes with this posterior belief on his ability. Denote the probability that an individual receives signal x_i given that he is of ability t_k by $p_k(x_i) = \Pr(x_i|t_k)$. Then the individuals posterior beliefs about his ability is given by

$$\Pr(t_k|x_i) = \frac{p_k(x_i) \cdot \frac{1}{5}}{\sum_j p_j(x_i) \cdot \frac{1}{5}} = \frac{p_k(x_i)}{\sum_j p_j(x_i)} \tag{1}$$

Now suppose that we ask the individual to predict the quintile in which his IQ score will fall, and promise him a payment if his prediction is correct. Clearly, an individual will pick the most likely quintile, i.e. the individual will indicate that he is *most likely* in ability quintile s_j where

$$s_j = \arg\max_j p_j(x_i) \tag{2}$$

Let us assume that our incentives are sufficient motivation for him to state the truth, and that he believes that our test is not biased. Then he will choose the most likely quintile, that is the quintile which maximizes the probability that his type is in that quintile. We call the theory that all subjects in the pool follow this procedure to determine their stated quintile the *Bayesian model*.

3.1 An example leading to overconfident judgments

Consider the following example: There are only two types, good and bad. The top two quintiles (40 percent) are good types, and the remaining three quintiles are bad types. This is the common prior. The only source of information the individuals receive is from a test that everybody takes. Good types pass the test for sure, while bad types only pass it with probability 50 percent. What is the probability that an individual is a good type if he passes the test?

$$\Pr(\text{good type}) = \frac{1 \cdot 0.4}{1 \cdot 0.4 + 0.5 \cdot 0.6} = \frac{4}{7} \tag{3}$$

Therefore, 70 percent of the population pass the test (all the good types, plus half the bad types). If asked by the experimenter whether they are more likely a good type or a bad type, therefore, they will all say that they think it's more likely that they are good types. Thus in this population, 70 percent say that they belong to the top 40 percent, much like we observe in our data presented in the previous section. Notice, however, that the beliefs are, on average correct. 70 percent of the population belief that they are good with probability

[3]Notice that we restrict attention to one draw from a signal structure, rather than, e.g., a dynamic acquisition of signals. We do this because dynamic acquisition signals can be redefined as a single draw from a meta-signal strucure.

4/7, and 30 percent belief that they are are good with probability 0. Thus, on average, the beliefs are correct.

Notice also that the feature of the test that lead to overconfidence in beliefs was that the test was easy. All good types, and so do some bad types. If the test were hard (all bad types fail, and so do half of the good types), the resulting judgments would be reversed and display underconfidence with only 20 percent believing that they are good types.

3.2 Testable Restrictions on Beliefs

Thus, incomplete information about one's abilities can lead to overconfident beliefs. However, the theory also imposes testable implications how the distribution of confidence judgments should be related to the true abilities of individuals. These are testable because the experimenter also observes the true score of the individual in the test. So he will have at the end of the experiment, for each of his subjects a pair of observations, (*true score, stated quintile).* The true score is not a precise measure of the IQ of an individual, of course. but that it is good enough so that we can ignore sampling error with respect to the quintiles.

Remember that since individuals are asked to pick the most likely quintile, therefore they will form their judgments using (1) and (2). Denote the expected fraction of individuals from true ability quintile k assigning themselves to quintile j based on the signal structure provided in (1) by $q_k(s_j)$. We call $q_k(s_j)$ the theoretical allocation function. It defines a 5-by-5 matrix of confidence judgments with the main diagonal containing the fractions that hold the correct beliefs about their abilities. Rows $q_k(s_j)$ with $k < j$ indicate individuals who hold overconfident beliefs, while rows with $k > j$ indicate the fraction of individuals holding underconfident beliefs. What restrictions does Bayesian updating place on this matrix? Because individuals pick the most likely quintile given the signal x_i that they received, the mode of individuals thinking they are in quintile k must actually have true ability quintile k. That is, Bayesian updating imposes that

$$q_k(s_k) = \max_l q_l(s_k) \tag{4}$$

In the appendix, we characterize this property more fully. We show that, in fact, the theoretical allocation function derived from equations (1) and (2) can be considered the canonical information structure that carries all the relevant information. Notice that we have so far assumed that all individuals draw signals from a common signal structure. This, however, is not a crucial assumption. If different individuals drew signals from different signal structures, this can be modeled as a meta signal-structure, in which individuals first observe from which sub-structure they will draw signals.

Thus, Bayesian updating implies a "main diagonal" condition if the theoretical allocation function is read as a matrix. But how can restrictions imposed by (4) be tested against the *empirical* allocation function $\hat{q}_k(s_j)$, i.e., the empirical distribution of confidence judgments as a function of the individuals' true ability? Intuitively, if there is strong evidence that the main diagonal condition is violated, this rejects Bayesian updating. Table 2 displays the empirical allocation function for the numeracy and IQ test. The table shows that in both cases, there is a clear violation of the main diagonal condition. For example, in the numeracy test, only 18 percent of the individuals from the third quintile put themselves into the third quintile. By contrast, 40 percent from the first quintile and 27 percent from the second quintile put themselves in the third quintile, in violation of (4). But is the violation significant? Since we don't know the underlying signal structure, how likely is it that a signal structure satisfying (4) generated the data in Table 2? We propose a test that gives the Bayesian model the best chance not to be rejected. We proceed the following way: We estimate the parameters of the theoretical signal structure by maximum likelihood subject to the constraint imposed by (4). That is, we estimate $q_k(s_j)$ to maximize

$$q_k(s_j) = \arg\max \sum_{j,k} n_{kj} q_k(s_j) \text{ subject to } q_k(s_k) = \max_l q_l(s_k) \tag{5}$$

where n_{kj} is the number of individuals of ability quintile k saying that they are in quintile j. This is a concave problem and maximization is straightforward. Denote the solution to (5) by $q_k^{ML}(s_j)$. Notice that this gives the best chance to the null hypothesis of Bayesian updating, since we pick q^{ML} as the one satisfying (4) that best fits the observed data. We then calculate the fit of q^{ML} to $\hat{q}$ as the mean square root error from each cell:

$$\hat{d} = \frac{1}{25}\sqrt{\sum_{j,k}(\hat{q}_k(s_j) - q_k^{ML}(s_j))^2} \tag{6}$$

The distance measure is $\hat{d}^{IQ} = 0.026$ for the IQ test, and $\hat{d}^{Num} = 0.033$ for the numeracy test. That is, the average deviation from the ML estimate of q is 2.6 percentage points in the IQ test and 3.3 percentage points in the Numeracy test. In order to assess whether the fit $\hat{d}$ is improbably bad, we generate 100,000 simulations of the same sample size as our data using q^{ML} as the data generating mechanism and calculate the distances d_n for each trial n. This provides us with an empirical distribution function for the distance measure d to calculate the probability that a draw from q^{ML} has a worse fit than the empirical allocation function $\hat{q}$. The p-values are $p = 0.005$ for the IQ test, and $p = 0.001$ for the numeracy test. Therefore, we clearly reject the hypothesis that our data is generated by imperfect information about ability and Bayesian updating.

4 Do Self-Image Concerns create Overconfidence?

The previous section tests and rejects a wide class of models that rely on Bayesian updating from a common prior after exogenous arrival of information. Other models have been developed to explain overconfidence arising endogenously as a function of individuals' choices.

Two recent papers (Kőszegi, 2006; Weinberg, 2006) have argued that a concern for self-image can lead to overconfidence. If individuals' utility depends on their belief about their ability, this can lead to an endogenous mechanism that produces results as if they were drawing signals from "easy test" signal structure in Benoît and Dubra (2007). This requires that utility is sufficiently "kinked" in the belief. Kőszegi (2006) provides an example in which an individual's utility discretely increases by v if the individual (rationally) believes that the chance that his ability t is below some threshold $\hat{t}$ is small. Formally, utility is given by

$$U(c, \hat{t}) = u(c) + v \cdot I(F(\hat{t}) \leq x) \tag{7}$$

where $F()$ is the c.d.f. of the individual's belief over his ability. To see how this can lead to overconfidence, assume that the individual's belief currently is that $F(\hat{t}) < x$ and that he is offered more information about his ability. The key mechanism generating overconfidence is that he will never seek more information in this case. More information only harbors the risk of erroneously revising his belief downward. Since there is no upside to seeking more information, the individual will decline. Conversely, if $F(\hat{t}) > x$, the individual will seek more information. If his belief is further revised downward, this leaves utility unchanged. If the individual receives a positive signal, he will gain utility v if $F'(\hat{t}) < x$ where $F'()$ is the c.d.f of beliefs incorporating the new information. Thus, this model can generate a pattern in which individuals with low beliefs will seek all the information they can find, while individuals with high beliefs will have less accurate information: Of all the individuals with initially low beliefs, all individuals with high ability will find out. By contrast, some of the individuals who initially had high beliefs will have received good signals by chance, but will not find out. The result is that too many individuals will belief they have high abilities.

We test the prediction of this model by testing the implication that individuals with high beliefs should be less likely to seek more information about their ability. Recall that after each test, we offered the subjects the opportunity to find out exactly how well they did relative to the others. We thus gave the individuals the chance to obtain more information, exactly as required in the model. This test also has the feature that it does not rely on the assumption of common priors. Rather, it measures the demand for information directly as a function of the individuals' beliefs.

Figure 4 displays the fraction of individuals information about their performance as a function of how well they did in the test, and their belief about their performance. Because of the small number of observation, we exclude individuals with beliefs in the bottom two quintiles. Panel A in Figure 4 displays the results for the IQ test, while the results for the Numeracy test are displayed in Panel B. Both Panel show a strong impact of the beliefs on the demand for information. However, in contrast to what is predicted by models in which the belief about ability enters the utility function, individuals with a higher belief are more likely to ask for the performance information. The figure also controls in a rudimentary way for differences in true abilities by splitting the sample into the top and bottom half of the performers. Thus, by comparing individuals with identical beliefs in the top and bottom half of the true abilities, we can gauge the impact of true ability on the demand for information. There is, essentially, no relationship between ability and the demand for information.

In order to formally test the model, we estimate the following probit equation

$$seek_i = \gamma q_i + x_i\beta + e_i \tag{8}$$

where *seek* is an indicator variable equal to 1 if the individuals seeks information about his performance in the test, and zero otherwise. We estimate the equation separately for the IQ and numeracy test. Our variable of interest is stated belief of the individual $q \in \{1, 2, ..., 5\}$ regarding the most likely quintile. The control variables x include controls for the test performance. We estimate a five-part linear spline in test performance, with the splines defined over quintiles in order to control for test performance in a flexible way. We also include personality characteristics as measured by the Minnesota Personality Questionnaire (MPQ, see **?**). They also include a large set of controls for socio-demographic differences across subjects: 5 dummy variables for education levels, 5 categories for ethnicities, a gender dummy, age and age squared, and household income.

The results are displayed in Tables 3 and 4 for the demand for information about one's performance in the IQ and numeracy test, respectively. The table displays marginal effects on the probability of seeking information, rather than the bare coefficient estimates. Both tables are structured the same way. In the first column, we test whether, as indicated by the figure, a higher belief increases the likelihood of demanding information. Column (1) in Table 3 controls for test performance using a flexible functional form. It shows that conditional on actual performafnce, the subject's belief about their performance predicts whether or not they seek information. More optimistic beliefs increase the likelihood of seeking information: a one-quintile increase in beliefs is associated with a 3 percentage point higher probability of demanding information about the test. The results are even stronger for the numeracy test, where a one-quintile increase in the belief leads to almost a 6 percentage point increase in the likelihood of seeking information. In both cases, the effects are statistically highly significant. Column (2) adds personality characteristics as controls, obtained from the MPQ. The only characteristics that is significant is harm avoidance. The effect is small, but lends itself to a plausible interpretation that individuals who are less risk averse are more likely to seek information. That is, individuals who feel less strongly about receiving bad news are more likely to seek information. In column (3), we add the socio-economic control variables. However, they have no effect on the coefficient of interest. Finally, in column (4), we also add the beliefs about the ability in the test as well as the beliefs about the ability in the other test as explanatory variables. Some individuals do change their evaluation over the course of the test (correlation between pre and post test beliefs: $\rho = 0.64$ for IQ and $\rho = 074$ for numeracy). Similarly, while beliefs are correlated across tests, they are not perfectly correlated ($rho = 0.54$ for beliefs after the test). This allows us to examine the specificity of the link between beliefs and the demand for information. Our results show that the link is highly specific. In Table 3, we see that only the most recent belief significantly correlated with the demand for information. Confidence in the numeracy test is uncorrelated with the demand for information, and so is confidence before the test, ceteris paribus. Our results are slightly weaker for numeracy, where we find a weak effect of confidence in IQ on the demand for information.

Overall these results clearly reject the driving force for overconfidence postulated by models of self-image concerns (Kőszegi, 2001; Weinberg, 2006). In fact, we find the opposite of what these models predict: More confident individuals are more likely to seek information. This is consistent with a model in which individuals consume the signals about their ability, not the resulting belief. However, this mechanism also tends to undermine overconfidence, as individuals with high confidence judgments are more likely to seek information, thus begging the question how overconfidence comes about in our subjects. One possibility is that individuals do not process information in a Bayesian manner. This interpretation is consistent with our evidence from the previous section, rejecting overconfidence as a consequence of incomplete information and learning. In particular, this explanation suggests that personality characteristics may be related to the misinterpretation of information. We explore this explanation in the next section.

5 Which Personality Traits create Overconfidence?

In this section, we examine whether differences in personality characteristics can explain differences in confidence judgments as implied by our interpretation of the results in the previous section. Our interpretation suggests that an individual's personality characteristics related to his desire for status and dominance, as well as his susceptibility to negative feedback should be influencing confidence judgments.

We focus on two dimensions that can readily be measured using personality scales: the first is an individual's desire to dominate others. The MPQ contains a dimension called "social potency," which measures exactly this trait. We predict that individuals who score high on social potency will have more confident beliefs, as they may interpret information in a more self-serving way. The MPQ also allows us to distinguish this from a more general desire to be connected to others, which is measured in a separate scale called "social closeness." It also allows us to distinguish the desire to dominate from general drive to achieve, using the "achievement" scale in the MPQ.

The second important dimension is how individuals respond to negative social feedback. We hypothesize that if individuals are worry-prone and feel vulnerable, this may moderate their beliefs about themselves to make it less likely to experience these negative social emotions. The MPQ also allows us to control for other aspects of risk preferences, such as a more general tendency towards prudence, as measured by the "harm avoidance" scale, and general pessimism captured by the "alienation" scale.

Figure 5 provides a first impression of the evidence. It shows confidence judgments and actual abilities for individuals with differing scores along personality traits. Each panel reflects a different personality trait. In each case, we cut the sample by the median trait score. For example, in Panel A, the first graph shows that about 30 percent of the individuals scoring below the median in social potency think they belong to the top 20 percent in the IQ distribution. By contrast, 55 percent of the individuals scoring above the median in social potency think they are in the top 20 percent. Each graph also contains the actual fraction of individuals scoring in the top 20 percent for each subsample. The graph shows virtually no difference between high- and low- social potency individuals in terms of actual ability. The results for confidence judgments in the numeracy test are very similar. Thus, social dominance appears to pick up quantitatively important differences in confidence judgments, while being unrelated to differences in actual abilities. Turning to the graph that cuts the sample by social closeness, we see no differences in confidence judgments. Thus, it is not the case that individuals who care more about sociability individuals are more confident in general, it is limited to the aspect of dominance relative to others. The third graph cuts the sample by the median of the stress reaction score. Individuals who are highly sensitive to social stress have substantially more timid judgments about their ability, as can be seen in the graph, while this is again not related to differences in actual abilities. Again, a very similar pattern emerges when we examine confidence judgments regarding the numeracy test in Panel B.

In order to examine these hypotheses using a formal statistical test, we estimate an ordered problit model of the form

$$q_i = k \text{ if } \alpha_{k-i} < \gamma MPQ_i + x_i\beta + e_i \leq \alpha_k \tag{9}$$

where MPQ is the full set of 11 dimensions of personality characteristics and x contains the same control variables as in the previous section. Tables 5 and 6 present the results. We report the marginal effects on the probability of believing that the individual thinks he is in the highest quintile.

In both tables, a range of personality characteristics are significant. Consistent with our interpretation, social potency is a highly significant predictor of confidence. A one-point increase in the scale leads to a 1.1 percent higher probability that the individual ranks himself in the top 20 percent. Given that the interquartile range on this index is 8, it predicts large and important differences in confidence judgments. Similarly, stress reaction predicts differences in confidence. Moving from the 25th to the 75th percentile in the stress reaction distribution (a 9-point increase) predicts a decrease in the highest confidence judgments by 8 percentage points. As we move to more restrictive specifications, using more flexible controls for cognitive ability and include our standard set of control variables, two personality characteristics remain significant: social potency strongly increases confidence judgments, while high scores in stress reactions reduces confidence. Thus, personality traits have a strong and significant impact on confidence judgments, in line with our interpretation that individuals interpret information in a biased way.

6 Conclusions

In this paper, we showed three results: First, we reject that overconfidence results from incomplete information about one's own ability and Bayesian learning (Benoît and Dubra, 2007). This implies that a standard economic model of unbiased information processing from a common prior cannot explain the evidence. Second, we rejected the hypothesis that optimistic beliefs about one's abilities lead individuals not to demand information about relative performance. We, therefore, reject the central prediction of models of self-image that can lead to overconfidence in beliefs (which we observe in our subjects). The opposite is true: We find a positive and highly significant association between more optimistic beliefs and demand for information about one's relative performance. This relationship is, as we have shown, specific to the belief about one's relative performance in test at hand, not the other test. Further, it is the belief after the test, not the belief about one's ability before the test that predicts the demand for information. Overall, we take this as evidence that it is the signal about good relative performance that the individual derives utility from. The high degree of specificity is consistent with this model, because it implies that individuals only seek out those signals that are most likely to be positive, thus maximizing the chance of a positive piece of news and a small utility kick.

References

Barber, B. and Odean, T. (2001), 'Boys Will be Boys: Gender, Overconfidence, and Common Stock Investment*', *Quarterly Journal of Economics* **116**(1), 261–292.

Bartling, B. and Fehr, E. a. (2009), 'Egalitarianism and competitiveness', *American Economic Review* **99**(2), 93 – 98.

Benoît, J.-P. and Dubra, J. (2007), Overconfidence?, MPRA Paper 6017, University Library of Munich, Germany.
URL: *http://ideas.repec.org/p/pra/mprapa/6017.html*

Burks, S., Carpenter, J., Goette, L., Monaco, K., Rustichini, A., Porter, K. and Hall, M. (2007), 'Using behavioral economic field experiments at a large motor carrier: The context and design of the truckers and turnover project', *NBER Working Paper* .

Burks, S., Carpenter, J., Goette, L. and Rustichini, A. (2009), 'Cognitive skills affect economic preferences, strategic behavior, and job attachment', *Proceedings of the National Academy of Sciences* **106**(19), 7745.

De Bondt, Werner, F. and Thaler, R. H. (1985), 'Does the stock market overreact', *Journal of Finance* **40**, 793–805.

Dohmen, T. J. and Falk, A. (2006), Performance pay and multi-dimensional sorting, IZA Discussion Papers 2001, Institute for the Study of Labor (IZA).

Kőszegi, B. (2006), 'Ego utility, overconfidence, and task choice', *Journal of the European Economic Association* **4**(4), 673–707.

Malmendier, U. and Tate, G. (2005), 'CEO overconfidence and corporate investment', *Journal of Finance* pp. 2661–2700.

Malmendier, U. and Tate, G. (2008), 'Who makes acquisitions? CEO overconfidence and the market's reaction', *Journal of Financial Economics* **89**(1), 20–43.

Merkle, C. and Weber, M. (2009), True overconfidence, Working paper, University of Mannheim.

Weinberg, B. (2006), A model of overconfidence, Working paper, Ohio State University.

A Tables and Figures

Figure 1: Distribution of Beliefs about One's Ability

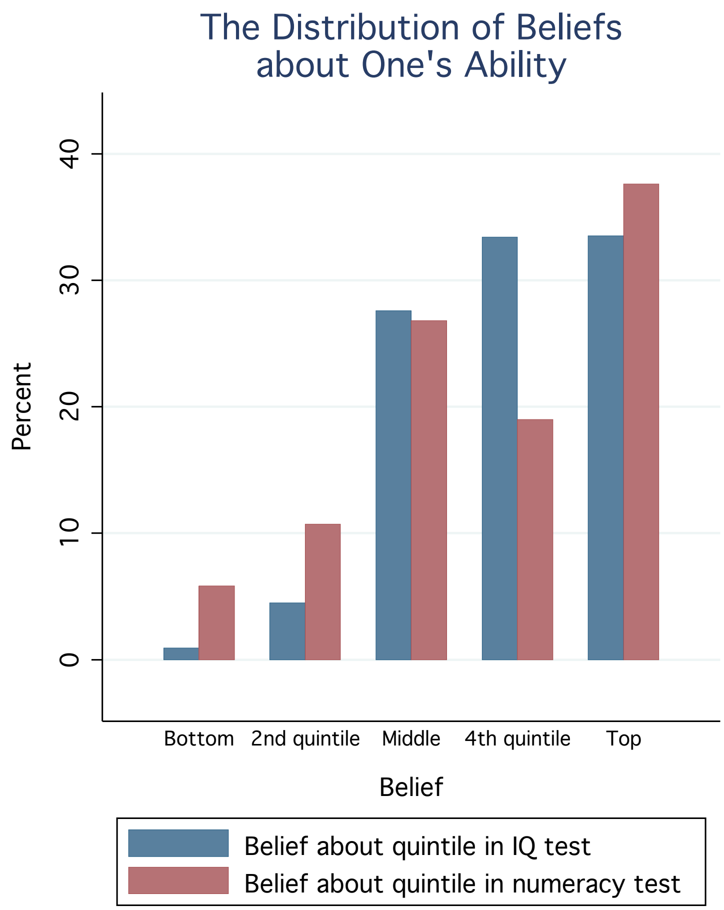

Figure 2: Confidence Judgments as a Function of Actual Ability: IQ.

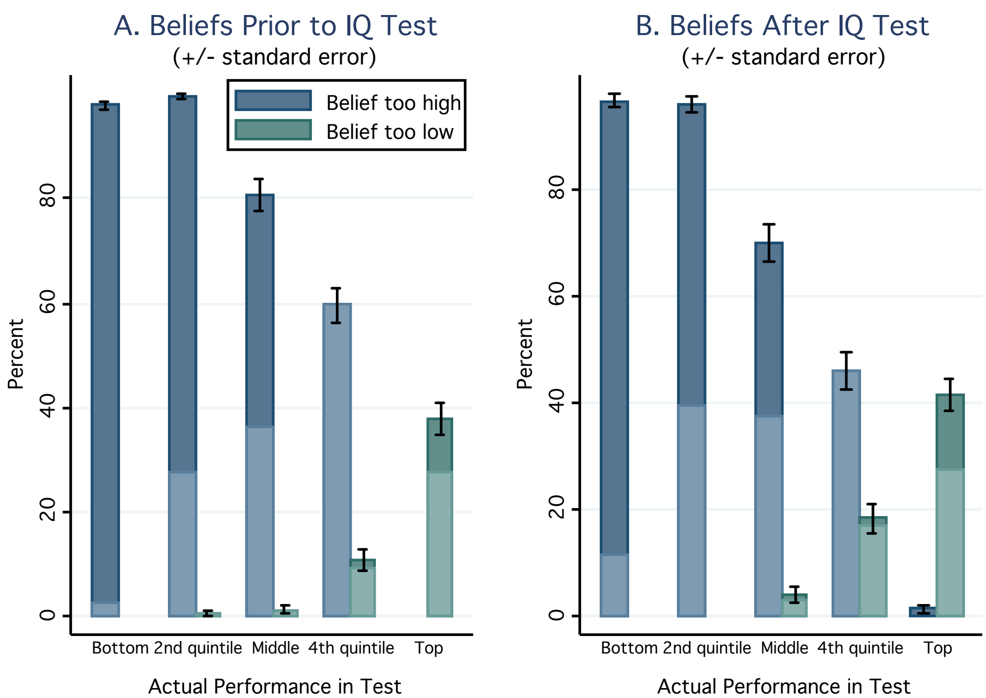

Figure 3: Confidence Judgments as a Function of Actual Ability: Numeracy.

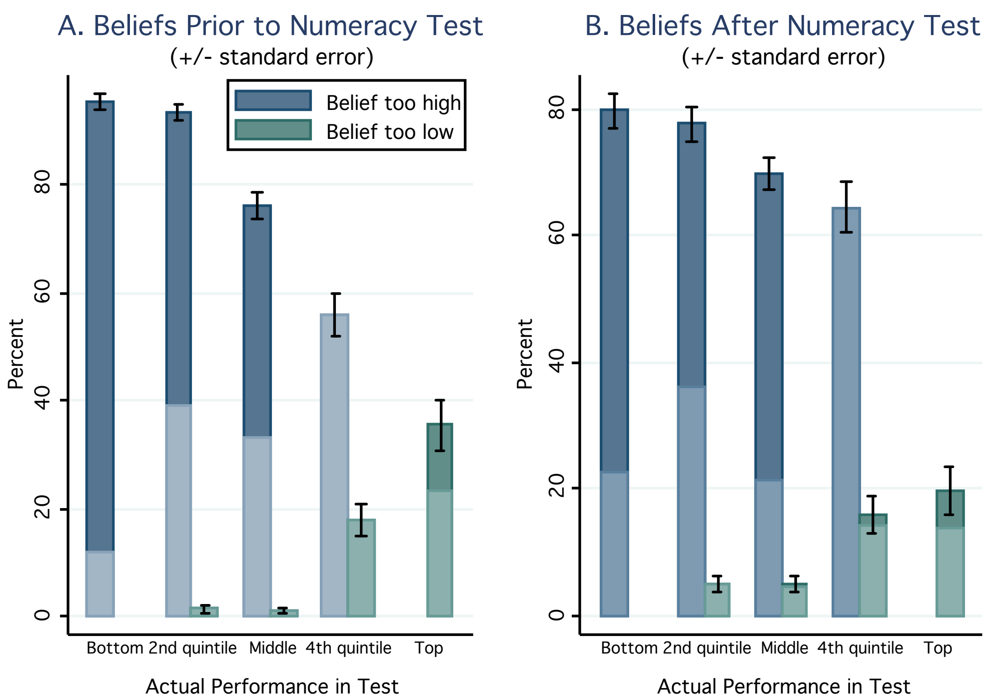

Figure 4: The demand for information

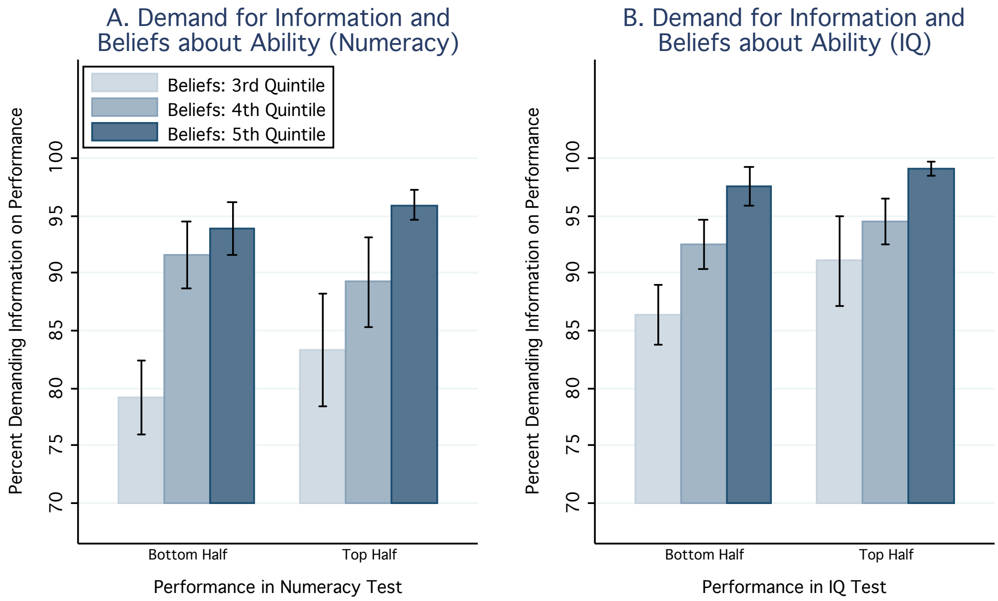

Notes: Caps indicate standard error of the respective mean.

Figure 5: Personality characteristics and confidence judgments

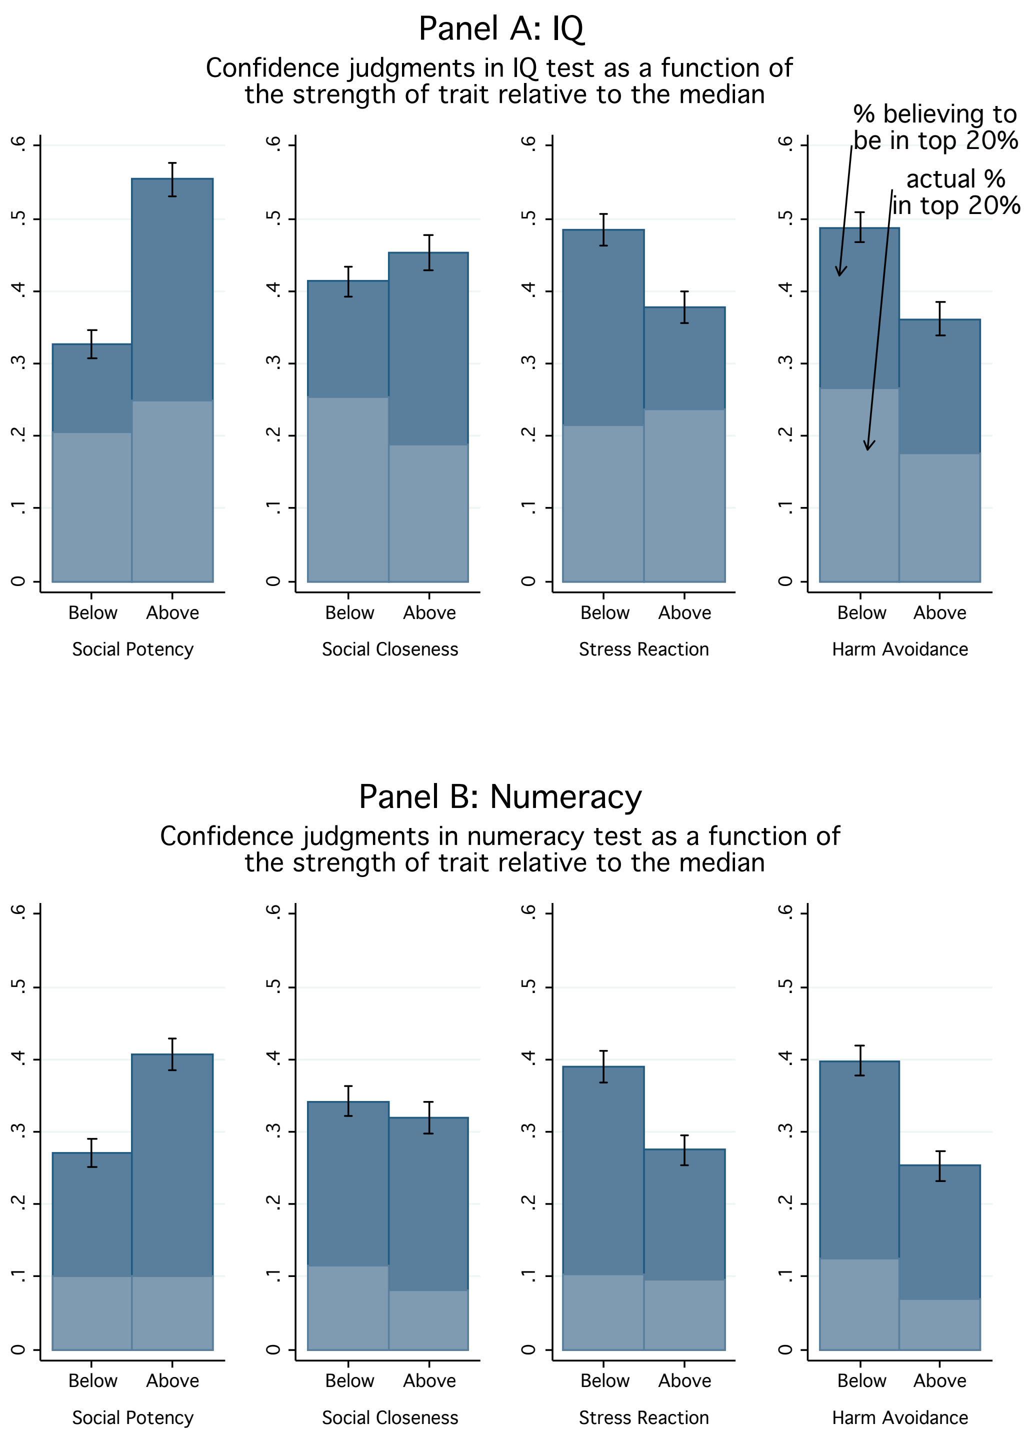

Notes: Caps indicate standard error of the respective mean.

Table 1: Descriptive Statistics

	Mean	Standard Deviation	Min	Max
Test Scores: Number of correct answers.				
Numeracy Test	8.54	2.55	1	12
IQ Test	45.33	8.15	1	60
Education: Highest level attained				
Middle School	3.9%			
High School	39.3%			
Technical School	14.9%			
Some College	33.2%			
College	6.5%			
Graduate School	2.3%			
Ethnic Categories:				
Caucasian	82.7%			
African-American	14.1%			
Indian	2.8%			
Asian	0.7%			
Latino	1.8%			
Other	1.0%			
Other Demographics:				
Age	37.43	10.90	21	69
Male	88.7%			
Household income (K)	52.66	27.07	10	150

Notes: $N = 888$ individuals.

Table 2: The Empirical Allocation functions $\hat{q}_k(s_j)$

Numeracy Test

	s_1	s_2	s_3	s_4	s_5
t_5	0.0	0.0	0.1	0.27	**0.62**
t_4	0.004	0.009	0.091	**0.298**	0.59
t_3	0.0	0.0125	**0.181**	0.362	0.443
t_2	0.004	**0.0**	0.272	0.377	0.345
t_1	**0.02**	0.02	0.401	0.376	0.175

IQ Test

	s_1	s_2	s_3	s_4	s_5
t_5	0.004	0.016	0.121	0.271	**0.579**
t_4	0.0	0.014	0.168	**0.355**	0.461
t_3	0.006	0.031	**0.262**	0.375	0.325
t_2	0.0	**0.04**	0.39	0.363	0.204
t_1	**0.033**	0.11	0.42	0.322	0.104

Notes: The empirical allocation function indicates for each ability quintile k, what fraction of individual put themselves in ability quintile j.

Table 3: The Demand for Information: IQ Test

Dependent Variable: Demand Information (=1)
Marginal Effects from Probit Estimates

	(1)	(2)	(3)	(4)
q_i^{IQ} after test	0.031***	0.029***	0.030***	0.030***
	(0.009)	(0.009)	(0.009)	(0.011)
q_i^{IQ} before test				– 0.004
				(0.011)
q_i^{NT} after test				0.005
				(0.007)
Piece-wise linear profile in test score				
first quintile	0.000	0.000	0.000	0.000
	(0.002)	(0.002)	(0.002)	(0.002)
second quintile	0.001	0.000	– 0.000	0.000
	(0.005)	(0.005)	(0.005)	(0.005)
third quintile	0.017	0.016	0.015	0.014
	(0.012)	(0.012)	(0.011)	(0.011)
fourth quintile	– 0.008	– 0.006	– 0.007	– 0.007
	(0.011)	(0.010)	(0.010)	(0.010)
fifth quintile	0.006	0.006	0.006	0.005
	(0.011)	(0.010)	(0.010)	(0.010)
Harm Avoidance		– 0.003**	– 0.003**	– 0.003**
		(0.002)	(0.002)	(0.002)
Social Closeness		0.002*	0.002	0.002
		(0.001)	(0.001)	(0.001)
Social Potency		– 0.001	– 0.001	– 0.001
		(0.002)	(0.002)	(0.002)
Stress Reaction		0.000	0.000	0.000
		(0.001)	(0.001)	(0.001)
Demographic controls?	No	No	Yes	Yes
p	0.000	0.001	0.003	0.005
N	838	838	826	825

Table 4: The Demand for Information: Numeracy Test

Dependent Variable: Demand Information (=1)
Marginal Effects from Probit Estimates

	(1)	(2)	(3)	(4)
q_i^{NT} after test	0.06***	0.057***	0.058***	0.039***
	(0.01)	(0.011)	(0.011)	(0.013)
q_i^{NT} before test				0.018
				(0.017)
q_i^{IQ} after test				0.028**
				(0.014)
Piece-wise linear profile in test score				
first quintile	0.022*	0.022*	0.023*	0.023*
	(0.014)	(0.013)	(0.014)	(0.013)
second quintile	0.011	0.001	0.002	0.002
	(0.020)	(0.020)	(0.020)	(0.019)
third quintile	0.038*	0.009	0.010	0.008
	(0.021)	(0.020)	(0.020)	(0.020)
fourth quintile	0.014	0.008	0.015	0.015
	(0.041)	(0.040)	(0.039)	(0.039)
fifth quintile	0.014	0.009	0.013	0.009
	(0.047)	(0.045)	(0.044)	(0.044)
Harm Avoidance		– 0.005**	– 0.005**	– 0.004*
		(0.002)	(0.002)	(0.002)
Social Closeness		0.003	0.002	0.002
		(0.002)	(0.002)	(0.002)
Social Potency		– 0.000	0.001	0.001
		(0.002)	(0.002)	(0.002)
Stress Reaction		0.002	0.002	0.002
		(0.002)	(0.002)	(0.002)
Demographic controls?	No	No	Yes	Yes
p	0.000	0.001	0.003	0.005
N	888	886	873	873

Table 5: Personality Characteristics and Confidence Judgments: IQ Test

Marginal Effects from Ordered Probit Model

	(1)	(2)	(3)	(4)
Absoption	0.004	0.002	0.001	
	(0.003)	(0.004)	(0.004)	
Achievement	0.010***	0.007*	0.007*	0.005
	(0.003)	(0.004)	(0.004)	(0.003)
Aggression	0.002	0.001	0.002	
	(0.003)	(0.003)	(0.003)	
Alienation	-0.007**	-0.003	-0.002	
	(0.003)	(0.003)	(0.003)	
Control	0.002	0.002	0.003	
	(0.003)	(0.004)	(0.004)	
Harm Avoidance	-0.009***	-0.008**	-0.007**	-0.007**
	(0.003)	(0.003)	(0.003)	(0.003)
Social Closeness	-0.004	-0.004	-0.001	-0.002
	(0.003)	(0.003)	(0.003)	(0.003)
Social Potency	0.014***	0.012***	0.012***	0.012***
	(0.003)	(0.003)	(0.003)	(0.003)
Stress Recation	-0.006**	-0.006**	-0.007**	-0.006***
	(0.003)	(0.003)	(0.003)	(0.002)
Traditionalism	-0.010***	-0.006*	-0.005	-0.005*
	(0.003)	(0.003)	(0.003)	(0.003)
Wellbeing	-0.005	-0.004	-0.004	
	(0.003)	(0.004)	(0.004)	
All 11 MPQ traits?	Yes	Yes	Yes	No
Control for performance	linear	linear	spline	spline
Demographics?	No	Yes	Yes	Yes
N	1063	1014	1014	1014

Table 6: Personality Characteristics and Confidence Judgments: Numeracy Test

Marginal Effects from Ordered Probit Model

	(1)	(2)	(3)	(4)
Absoption	0.006**	0.006*	0.004	
	(0.003)	(0.003)	(0.003)	
Achievement	0.004	0.001	0.001	0.000
	(0.003)	(0.003)	(0.003)	(0.003)
Aggression	-0.001	-0.001	0.000	
	(0.003)	(0.003)	(0.003)	
Alienation	-0.006**	-0.002	0.000	
	(0.003)	(0.003)	(0.003)	
Control	-0.001	-0.001	0.002	
	(0.003)	(0.003)	(0.003)	
Harm Avoidance	-0.008***	-0.004	-0.004	-0.004
	(0.003)	(0.003)	(0.003)	(0.003)
Social Closeness	-0.006**	-0.006**	-0.003	-0.004
	(0.003)	(0.003)	(0.003)	(0.002)
Social Potency	0.011***	0.008***	0.007**	0.007***
	(0.003)	(0.003)	(0.003)	(0.003)
Stress Recation	-0.007***	-0.008***	-0.009***	-0.007***
	(0.002)	(0.003)	(0.003)	(0.002)
Traditionalism	-0.007**	-0.003	-0.002	-0.002
	(0.003)	(0.003)	(0.003)	(0.003)
Wellbeing	-0.006*	-0.005*	-0.005	
	(0.003)	(0.003)	(0.003)	
All 11 MPQ traits?	Yes	Yes	Yes	No
Control for performance	linear	linear	spline	spline
Demographics?	No	Yes	Yes	Yes
N	1063	1014	1014	1014